DNS Operations
Internet-Draft
Expires: January 19, 2006

K. Fujiwara
JPRS
K. Ishibashi
K. Toyama
NTT PF Labs
C. Yoshimura
NTT Communications

July 18, 2005

DNS authoritative server misconfiguration and countermeasure in resolver

draft-fujiwara-dnsop-bad-dns-auth-03.txt

Status of this Memo

Copyright Notice

Abstract

This memo describes misconfigurations of DNS authoritative servers and its effect to old DNS iterative resolver servers we experienced. We recommend re-checking DNS authoritative server configuration and advise using newer iterative resolver server implementations. The

recommendations made in this document are based on analysis of abnormal DNS resolver server load at large ISP resolver server which has many customers. This is not protocol issue.

Table of Contents

1. Introduction

There are many misconfigured DNS authoritative servers. They have large RRSets whose response size exceeds 512 octets, they does not support EDNS0 extension, and they does not answer TCP DNS query. In this case, over 512 octets RRSets cannot be resolved.

This memo describes that combination of misconfigurations at authoritative servers can create significant overloads on resolver servers, especially old but spreaded BIND 8.

While there are reports on the observations of query traffic to root or top-level domain servers and the recommendations to the resolver servers to reduce anomalies on the servers [draft-ietf-dnsop-bad-dns-res-03], [WESSELS04], this memo intends to notify to the operators of authoritative servers that their configuration can lead resolver servers' problems.

In the following sections, we provide a detailed explanation of the problem. We then recommend to re-check the configurations of authoritative servers to avoid the problem. At last, we describe iterative resolver server's recommendation.

2. Problem Description

DNS message size is limited to 512 octets in original UDP packet [RFC1035]. However, it is possible to write large RRSet which exceeds 512 octets. A typical case observed is a response with PTR RRSet for an IP address which is assigned for many (over 300) domain names [TOYAMA04]. Another case, many A RRs to one domain name for load balancing or by writing many SRV RRs to large domain name for Active Directory.

Iterative resolver servers send queries to authoritative servers. If one authoritative server which returns such large response does not support EDNS0 and sufficient maximum payload size [RFC2671], the server returns truncated response (TC bit = 1) to the resolver server. Then the resolver server tries to get whole message by using TCP transport.

If the authoritative server which returns such large response does not support TCP transport or filters TCP DNS port, the DNS query fails. There are multiple authoritative servers for the record, the resolver server repeats the sequence for all the authoritative servers. If all authoritative servers for the RRSet are misconfigured and does not answer by TCP, the DNS query cannot be resolved and the resolver server responds ServFail error to stub resolver. If the authoritative server filters TCP DNS port and does

not send TCP reset, the resolver server must wait till TCP timeout.

This RRSet status cannot be cached by both resolver servers and stub resolvers. This case corresponds to [RFC2308] section 7.2 Dead / Unreachable Server. Dead or unreachable server information may be cached in 5 minutes. As the result, there are many queries to misconfigured authoritative servers.

The problem we faced is significant overloads of BIND 8 resolver servers. BIND 8 resolver server starts a iterative query at every query from stub resolvers (when there are no cached data for that query) and it keeps a TCP SYN_SENT state for some interval. There were many queries for unresolvable RRSet and keeping many TCP states increased the load of the resolver server. This phenomenon impacted significantly the resolver server performance. These unresolvable RRSets are well-known addresses.

While BIND 9 iterative resolver server is resolving one domainname, it does not try to resolve the same queries and it will answer same response at the first query completion. Therefore, BIND 9 is not affected.

3. Re-checking of Authoritative servers

Authoritative DNS servers with large RRSets whose response size may exceeds 512 octets must answer TCP DNS query and should support EDNS0. Or the RRset cannt be resolved.

System administrators manage TCP filters carefully and some of them does not know about DNS. As a result, some administrators filters their DNS server's TCP port.

Therefore, the operators of the authoritative servers should know about DNS and should re-check the configuration of their servers.

4. Iterative resolver server requirements

There are authoritative DNS servers with TCP filer problems. ISP DNS resolver servers must resolve or answer any query which ISP customer queries. Even if they receive unresolvable queries, they must work well. So, using tough iterative resolver server implementation is necessary.

Currently, any BIND 8 version have this weak-point. Using BIND 9 is one solution.

5. Conclusion

Reducing unresolvable RRSets is necessary. But there still exist misconfigurations. Iterative resolver servers which support many users must be tough. So, using older implementation should be deprecated.

6. Security considerations

Older iterative resolver server implementations especially old but spreaded BIND 8 may have weak-points. Using older and weak implementations should be deprecated.

7. References

[TOYAMA04]
Toyama, K., "DNS Anomalies and Their Impact on DNS Cache Servers", NANOG 32, October 2004.

[draft-ietf-dnsop-respsize-01]
Vixie, P. and A. Kato, "DNS Response Size Issues (work in progress)", July 2004.

[draft-ietf-dnsop-bad-dns-res-03]
Larson, M. and P. Barber, "Observed DNS Resolution Misbehavior", October 2004.

[WESSELS04]
Wessels, D., "Is Your Caching Resolver Polluting the Internet?", SIGCOMM Network Troubleshooting, August 2004.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, March 1998.

[RFC2671] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, August 1999.

Authors' Addresses

Kazunori Fujiwara
Japan Registry Services Co.,Ltd.
Chiyoda First Bldg. East 13F
3-8-1 Nishi-Kanda Chiyoda-ku
Tokyo 101-0065
JAPAN

Phone: +81-3-5215-8451
Email: fujiwara@jprs.co.jp

Keisuke Ishibashi
Information Sharing Platform Laboratories, Nippon Telegraph and Telephone Corporation
3-9-11 Midori-cho
Musashino-shi
Tokyo 180-8585
JAPAN

Phone: +81-422-59-3407
Email: ishibashi.keisuke@lab.ntt.co.jp

Katsuyasu Toyama
Information Sharing Platform Laboratories, Nippon Telegraph and Telephone Corporation
3-9-11 Midori-cho
Musashino-shi
Tokyo 180-8585
JAPAN

Phone: +81-422-59-7906
Email: toyama.katsuyasu@lab.ntt.co.jp

Chika Yoshimura
NTT Communications Corporation
NTT OTEMACHI BLDG.
2-3-5 Otemachi, Chiyoda-ku
Tokyo 100-0004
JAPAN

Phone: +81-3-6800-6113
Email: yosimura@ocn.ad.jp

Appendix A. Acknowledgements

We would like to thank Ichiro Mizukoshi, Haruhiko Ohshima, Masahiro

Ishino, Chika Yoshimura, Tsuyoshi Toyono, Hirotaka Matsuoka, Yasuhiro Morisita, and Bill Manning.

Intellectual Property Statement

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.