

Network Working Group
Request for Comments: 2870
Obsoletes: 2010
BCP: 40
Category: Best Current Practice

R. Bush
Verio
D. Karrenberg
RIPE NCC
M. Koster
Network Solutions
R. Plzak
SAIC
June 2000

Root Name Server Operational Requirements

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

As the internet becomes increasingly critical to the world's social and economic infrastructure, attention has rightly focused on the correct, safe, reliable, and secure operation of the internet infrastructure itself. The root domain name servers are seen as a crucial part of that technical infrastructure. The primary focus of this document is to provide guidelines for operation of the root name servers. Other major zone server operators (gTLDs, ccTLDs, major zones) may also find it useful. These guidelines are intended to meet the perceived societal needs without overly prescribing technical details.

1. Background

The resolution of domain names on the internet is critically dependent on the proper, safe, and secure operation of the root domain name servers. Currently, these dozen or so servers are provided and operated by a very competent and trusted group of volunteers. This document does not propose to change that, but merely to provide formal guidelines so that the community understands how and why this is done.

- 1.1 The Internet Corporation for Assigned Names and Numbers (ICANN) has become responsible for the operation of the root servers. The ICANN has appointed a Root Server System Advisory Committee (RSSAC) to give technical and operational advice to the ICANN board. The ICANN and the RSSAC look to the IETF to provide engineering standards.
- 1.2 The root servers serve the root, aka ".", zone. Although today some of the root servers also serve some TLDs (top level domains) such as gTLDs (COM, NET, ORG, etc.), infrastructural TLDs such as INT and IN-ADDR.ARPA, and some ccTLDs (country code TLDs, e.g. SE for Sweden), this is likely to change (see 2.5).
- 1.3 The root servers are neither involved with nor dependent upon the 'whois' data.
- 1.4 The domain name system has proven to be sufficiently robust that we are confident that the, presumably temporary, loss of most of the root servers should not significantly affect operation of the internet.
- 1.5 Experience has shown that the internet is quite vulnerable to incorrect data in the root zone or TLDs. Hence authentication, validation, and security of these data are of great concern.

2. The Servers Themselves

The following are requirements for the technical details of the root servers themselves:

- 2.1 It would be short-sighted of this document to specify particular hardware, operating systems, or name serving software. Variations in these areas would actually add overall robustness.
- 2.2 Each server MUST run software which correctly implements the IETF standards for the DNS, currently [RFC1035] [RFC2181]. While there are no formal test suites for standards compliance, the maintainers of software used on root servers are expected to take all reasonable actions to conform to the IETF's then current documented expectations.
- 2.3 At any time, each server MUST be able to handle a load of requests for root data which is three times the measured peak of such requests on the most loaded server in then current normal conditions. This is usually expressed in requests per second. This is intended to ensure continued operation of root services should two thirds of the servers be taken out of operation, whether by intent, accident, or malice.

- 2.4 Each root server should have sufficient connectivity to the internet to support the bandwidth needs of the above requirement. Connectivity to the internet SHOULD be as diverse as possible.

Root servers SHOULD have mechanisms in place to accept IP connectivity to the root server from any internet provider delivering connectivity at their own cost.

- 2.5 Servers MUST provide authoritative responses only from the zones they serve. The servers MUST disable recursive lookup, forwarding, or any other function that may allow them to provide cached answers. They also MUST NOT provide secondary service for any zones other than the root and root-servers.net zones. These restrictions help prevent undue load on the root servers and reduce the chance of their caching incorrect data.

- 2.6 Root servers MUST answer queries from any internet host, i.e. may not block root name resolution from any valid IP address, except in the case of queries causing operational problems, in which case the blocking SHOULD last only as long as the problem, and be as specific as reasonably possible.

- 2.7 Root servers SHOULD NOT answer AXFR, or other zone transfer, queries from clients other than other root servers. This restriction is intended to, among other things, prevent unnecessary load on the root servers as advice has been heard such as "To avoid having a corruptible cache, make your server a stealth secondary for the root zone." The root servers MAY put the root zone up for ftp or other access on one or more less critical servers.

- 2.8 Servers MUST generate checksums when sending UDP datagrams and MUST verify checksums when receiving UDP datagrams containing a non-zero checksum.

3. Security Considerations

The servers need both physical and protocol security as well as unambiguous authentication of their responses.

- 3.1 Physical security MUST be ensured in a manner expected of data centers critical to a major enterprise.

- 3.1.1 Whether or not the overall site in which a root server is located has access control, the specific area in which the root server is located MUST have positive access control, i.e. the number of individuals permitted access to the area MUST be limited, controlled, and recorded. At a

minimum, control measures SHOULD be either mechanical or electronic locks. Physical security MAY be enhanced by the use of intrusion detection and motion sensors, multiple serial access points, security personnel, etc.

3.1.2 Unless there is documentable experience that the local power grid is more reliable than the MTBF of a UPS (i.e. five to ten years), power continuity for at least 48 hours MUST be assured, whether through on-site batteries, on-site power generation, or some combination thereof. This MUST supply the server itself, as well as the infrastructure necessary to connect the server to the internet. There MUST be procedures which ensure that power fallback mechanisms and supplies are tested no less frequently than the specifications and recommendations of the manufacturer.

3.1.3 Fire detection and/or retardation MUST be provided.

3.1.4 Provision MUST be made for rapid return to operation after a system outage. This SHOULD involve backup of systems software and configuration. But SHOULD also involve backup hardware which is pre-configured and ready to take over operation, which MAY require manual procedures.

3.2 Network security should be of the level provided for critical infrastructure of a major commercial enterprise.

3.2.1 The root servers themselves MUST NOT provide services other than root name service e.g. remote internet protocols such as http, telnet, rlogin, ftp, etc. The only login accounts permitted should be for the server administrator(s). "Root" or "privileged user" access MUST NOT be permitted except through an intermediate user account.

Servers MUST have a secure mechanism for remote administrative access and maintenance. Failures happen; given the 24x7 support requirement (per 4.5), there will be times when something breaks badly enough that senior wizards will have to connect remotely. Remote logins MUST be protected by a secure means that is strongly authenticated and encrypted, and sites from which remote login is allowed MUST be protected and hardened.

3.2.2 Root name servers SHOULD NOT trust other hosts, except secondary servers trusting the primary server, for matters of authentication, encryption keys, or other access or

security information. If a root operator uses kerberos authentication to manage access to the root server, then the associated kerberos key server **MUST** be protected with the same prudence as the root server itself. This applies to all related services which are trusted in any manner.

- 3.2.3 The LAN segment(s) on which a root server is homed **MUST NOT** also home crackable hosts. I.e. the LAN segments should be switched or routed so there is no possibility of masquerading. Some LAN switches aren't suitable for security purposes, there have been published attacks on their filtering. While these can often be prevented by careful configuration, extreme prudence is recommended. It is best if the LAN segment simply does not have any other hosts on it.
- 3.2.4 The LAN segment(s) on which a root server is homed **SHOULD** be separately firewalled or packet filtered to discourage network access to any port other than those needed for name service.
- 3.2.5 The root servers **SHOULD** have their clocks synchronized via NTP [RFC1305] [RFC2030] or similar mechanisms, in as secure manner as possible. For this purpose, servers and their associated firewalls **SHOULD** allow the root servers to be NTP clients. Root servers **MUST NOT** act as NTP peers or servers.
- 3.2.6 All attempts at intrusion or other compromise **SHOULD** be logged, and all such logs from all root servers **SHOULD** be analyzed by a cooperative security team communicating with all server operators to look for patterns, serious attempts, etc. Servers **SHOULD** log in GMT to facilitate log comparison.
- 3.2.7 Server logging **SHOULD** be to separate hosts which **SHOULD** be protected similarly to the root servers themselves.
- 3.2.8 The server **SHOULD** be protected from attacks based on source routing. The server **MUST NOT** rely on address- or name-based authentication.
- 3.2.9 The network on which the server is homed **SHOULD** have in-addr.arpa service.

- 3.3 Protocol authentication and security are required to ensure that data presented by the root servers are those created by those authorized to maintain the root zone data.

- 3.3.1 The root zone MUST be signed by the Internet Assigned Numbers Authority (IANA) in accordance with DNSSEC, see [RFC2535] or its replacements. It is understood that DNSSEC is not yet deployable on some common platforms, but will be deployed when supported.
- 3.3.2 Root servers MUST be DNSSEC-capable so that queries may be authenticated by clients with security and authentication concerns. It is understood that DNSSEC is not yet deployable on some common platforms, but will be deployed when supported.
- 3.3.3 Transfer of the root zone between root servers MUST be authenticated and be as secure as reasonably possible. Out of band security validation of updates MUST be supported. Servers MUST use DNSSEC to authenticate root zones received from other servers. It is understood that DNSSEC is not yet deployable on some common platforms, but will be deployed when supported.
- 3.3.4 A 'hidden primary' server, which only allows access by the authorized secondary root servers, MAY be used.
- 3.3.5 Root zone updates SHOULD only progress after a number of heuristic checks designed to detect erroneous updates have been passed. In case the update fails the tests, human intervention MUST be requested.
- 3.3.6 Root zone updates SHOULD normally be effective no later than 6 hours from notification of the root server operator.
- 3.3.7 A special procedure for emergency updates SHOULD be defined. Updates initiated by the emergency procedure SHOULD be made no later than 12 hours after notification.
- 3.3.8 In the advent of a critical network failure, each root server MUST have a method to update the root zone data via a medium which is delivered through an alternative, non-network, path.
- 3.3.9 Each root MUST keep global statistics on the amount and types of queries received/answered on a daily basis. These statistics must be made available to RSSAC and RSSAC sponsored researchers to help determine how to better deploy these machines more efficiently across the

internet. Each root MAY collect data snapshots to help determine data points such as DNS query storms, significant implementation bugs, etc.

4. Communications

Communications and coordination between root server operators and between the operators and the IANA and ICANN are necessary.

- 4.1 Planned outages and other down times SHOULD be coordinated between root server operators to ensure that a significant number of the root servers are not all down at the same time. Preannouncement of planned outages also keeps other operators from wasting time wondering about any anomalies.
- 4.2 Root server operators SHOULD coordinate backup timing so that many servers are not off-line being backed up at the same time. Backups SHOULD be frequently transferred off site.
- 4.3 Root server operators SHOULD exchange log files, particularly as they relate to security, loading, and other significant events. This MAY be through a central log coordination point, or MAY be informal.
- 4.4 Statistics as they concern usage rates, loading, and resource utilization SHOULD be exchanged between operators, and MUST be reported to the IANA for planning and reporting purposes.
- 4.5 Root name server administrative personnel MUST be available to provide service 24 hours a day, 7 days per week. On call personnel MAY be used to provide this service outside of normal working hours.

5. Acknowledgements

The authors would like to thank Scott Bradner, Robert Elz, Chris Fletcher, John Klensin, Steve Bellovin, and Vern Paxson for their constructive comments.

6. References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.
- [RFC2030] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 2030, October 1996.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.
- [RFC2535] Eastlake, D. and C. Kaufman, "Domain Name System Security Extensions", RFC 2535, March 1999.

7. Authors' Addresses

Randy Bush
Verio, Inc.
5147 Crystal Springs
Bainbridge Island, WA US-98110

Phone: +1 206 780 0431
EMail: randy@psg.com

Daniel Karrenberg
RIPE Network Coordination Centre (NCC)
Singel 258
NL-1016 AB Amsterdam
Netherlands

Phone: +31 20 535 4444
EMail: daniel.karrenberg@ripe.net

Mark Kusters
Network Solutions
505 Huntmar Park Drive
Herndon, VA 22070-5100

Phone: +1 703 742 0400
EMail: markk@netsol.com

Raymond Plzak
SAIC
1710 Goodridge Drive
McLean, Virginia 22102
+1 703 821 6535

EMail: plzakr@saic.com

8. Specification of Requirements

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.

9. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

