DNSOP
Internet-Draft
Expires: December 3, 2005

K. Lindqvist
Netnod Internet Exchange

June 2005

# Operation of TLD zones
# draft-kurtis-tld-ops-00.txt

Status of this Memo



Copyright Notice



Abstract

The Internet is today the defacto standard packet network for a lot of critical communications. The Internet in turn have a heavy dependency on the Domain Name System (DNS) for it's "normal" operations. The IETF in June 2000 described the operating requirements [1]for the so called root-servers that defines the root of the DNS lookup tree. Similar requirements could where deemed needed be applied to DNS infrastructure at other levels of the DNS tree as well. This document analyses these requirements and what can be done to ensure a reliable DNS infrastructure.





Table of Contents

## 1. Introduction

The Internet is today the defacto standard packet network for a lot of critical communications. The Internet in turn have a heavy dependency on the Domain Name System (DNS) for it's "normal" operations. The IETF in June 2000 described the operating requirements [1]for the so called root-servers that defines the root of the DNS lookup tree. So far though, operational requirements of the second tier of the DNS has not been defined. This document tries to analyze and define the operational requirements for the second tier in the DNS lookup hierarchy.

The second tier infrastructure in the DNS is called the Top Level Domains (TLDs). These include both generic TLDs (gTLDs) such as .com and .info and country code TLDs (ccTLDs). In the hierarchy immediately under the TLDs, we will have either directly registered domain names in use by an entity for lookup of their services and host names. Alternatively we will find a third their hierarchy grouping the same type of names, for example edu.uk for educational institutions in the UK, and co.uk for commercial companies in the UK. The latter method is mostly used to provide better scaling for the TLD. In this case this document would be applicable to both the second and the third tier.

The service level experience of the end-user with regard to DNS is

dependent on several levels of hierarchy. Usually each level will be the responsibility of multiple entities, such as the root-server operators, the TLD slave-server operators and the operator of the actual leaf zone. The weakest operation in that chain can make the service unreachable and result in a failure for an end node trying to access a particular service address, or host name. The root-server operations are as noted earlier described in an IETF/RFC. This document tries to outline the same requirements for the TLDs. This leaves the operator of the leaf zone. This document should however also provide useful information for operators of leaf zones, but some of the requirements may be considered to strict for most leaf zones.

## 2. Terminology

Service Address: an IP address associated with a particular service (e.g. the address of a nameserver).

## 3. Use of this document

While this document provides for a set of operational requirements for a TLD, they are not to be taken as absolute requirements. Needs and capabilities of TLDs will vary over the different operational aspects in this document. Some of them are still to be considered as





minimal requirements, the slave-server operations, registry operations and protocol requirements.

This document is not to been seen as an absolute requirements document, but rather as a check-list for TLD registrars, registries and TLD slave-server operators. The intent is also no to try and "score" the various requirements. The priority among them will vary very much between different TLDs and uses.

## 4. Analysis of requirements

The requirements in this document are targeting maximum operational stability, security and resilience for a TLD operation. This reflects the critical dependency many systems today have on the Internet, and the role the DNS has come to play. The requirements are not targeting business practices, or take into account affordability or need. This analysis is left to the users and readers of this document.

## 5. Operational requirements

### 5.1 Slave server operations

#### 5.1.1 Location

The ideal location of slave-servers can be broken down into two sub-categories. The ideal physical location, which is discussed in Section 5.1.2, and location in the network topology.

Finding the ideal location for the TLD slave-servers in the topology is a complex analysis. Factors that should be taken into account are

- Ensure optimal performance for the target community (such as a particular country for a ccTLD). I.e try to be as topologically close to the largest user base as possible.
- The slave servers should be as topologically diverse as is feasible to avoid a infrastructure problem taking out one or more

of the servers.
- Locating slave-servers where they will have the best and easiest access to a large number of networks will provide better performance and give more networks shorter and more direct access to the data.

#### 5.1.2 Physical requirements

The location of the slave-servers should meet co-location industry standards with regard to the physical environment. This includes





Redundant power feeding: The servers should be located so that a single interrupt of power into the site where these are located does not effect the servers, or the network infrastructure servicing the slave-servers.
Cooling: The site should provide adequate and redundant cooling that also operates adequately in the event of a power failure.
Access: The physical location should have restricted access, and only allow required operational staff to gain physical and control of the slave-servers.
Fire supresssion: The physical location should provide adequate fire fighting and/or suppression equipment.

#### 5.1.3 Software requirements

The slave-servers should be held to normal industry standards. This includes standard operational practices as

Patches: The server operating system and software should be kept uptodate with the most recent security patches for the software.
Backups: The slave-server systems should be backed up regularly. This includes zone data, system configurations and other data and configurations needed for a quick system recovery in the case of a failure. The procedures and systems for backup should be tested for resteration regularly.
Capacity: The system software should be otpimally tuned to serve the requirements of the zone in question. The system should be able to handle at least three times the normal load of questions per seconds.
Remote management: The server system should allow for secure, encrypted and authenticated remote management.
Time syncronization: Servers should for accuracy of logging have their clocks syncronized using the Network Time Protocol [2].

#### 5.1.4 Protocol requirements

The slave-servers must be running software that supports the current set of DNS protocol standards (as of writing RFC1034 [3],RFC1035 [4],RFC2181 [5]. For TLDs that are claiming Secure DNS support, the server software also should support RFC4033 [6],RFC4034 [7], RFC4035 [8].

In addition to supporting the above protocol standards there are also a number of configuration parameters the slave-service should follow. These are

Recursion: The slave-service should not be providing recursive name-service.

In-Bailiwick-glue: The zone should be configured to use "In-Bailiwick-glue", as that will give the TLD operator full control over the entire delegation chain and give the TLD operational control over the stability of the service.

Limited service The TLD slave-service should not at the same time be authorative for relatives (lower tiers) to the TLD zone. This avoids collapsing signed delegation data when DNSSEC is used.

### 5.2 Slave server infrastructure

A critical factor in guaranteeing the stability, resilience and redundancy of the TLD slave-service is the network infrastructure that connects the physical servers. With the increase in Distributed Denial of Service(DDoS) attacks, this is also an increasing area of concerns for TLD operators. Configuring and dimensioning the network infrastrucutre is therefor becoming inreasingly important for the TLD operator.

#### 5.2.1 Dimensioning the infrastructure

The network infrastructure should be dimensioned to handle a network load of three times the normal load, measured in packets-per-seconds, pps. The infrastructure should also be able of dropping packets at line-rate in order to protect the service in the case of an attack. What "line-rate" corresponds to in absolute numbers is depending on local conditions, such as affordability, need and uplink speeds.

#### 5.2.2 Following common practice

The network infrastructure fronting the slave-service should be configured accoring to normal industry practice. This includes at a minimum

Network separation: The network connection should not be shared with other, non-trusted hosts. In other words, the connection should be provided on a switched or routed infrastructure to avoid wiretapping and/or spoofing of packets.

Ingress filtering: The network infrastructure should filter out spoofed packets using the methods descibed in RFC3704/BCP84 [9].

Network protection: The infrastructure should apply some form of packet filtering methods that only allows traffic to ports needed for the service and management traffic from well-known management networks if needed.



Network security: The network infrastructure should be configured accoring to industry standard, for example as described in FILL_IN. The network infrastructure should also onyl be remotly and physically accessiable by operational staff. Remote access must be authenticated and encrypted.

### 5.3 Registry operations

#### 5.3.1 Registry/Registrar interface

FILL_IN perhaps some text on EPP etc.

## 6. Security Considerations

## 7. Protocol Considerations

This document does not impose any protocol considerations.

## 8. IANA Considerations

This document requests no action from IANA.

## 9. Acknowledgements

The following people have contributed to this document: David Meyer, Peter Koch, Johan Ihren, Patrik Falstrom. The author would like to extend great appreciation to them for the reviews, and text they have contributed with.

## 10. References

[1] Bush, R., Karrenberg, D., Kosters, M., and R. Plzak, "Root Name Server Operational Requirements", BCP 40, RFC 2870, June 2000.

[2] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.

[3] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

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

[5] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.





[6] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[7] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[8] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

[9] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

## Author's Address

Kurt Erik Lindqvist
Netnod Internet Exchange

Bellmansgatan 30
118 47 Stockholm
Sweden

Email: kurtis@kurtis.pp.se
URI: http://www.netnod.se/





## Intellectual Property Statement



## Disclaimer of Validity

Acknowledgment

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