DNS Infrastructure Recommendation

Of the

Security and Stability Advisory Committee

SAC 005

Document 005 Version 1

1 November 2003

Table Of Contents

Executive Summary

A key element of the DNS infrastructure is the delegation of zones. Beginning with the root of the DNS ("."), each zone administrator has the authority to delegate sub-zones to other responsible parties. Each sub-zone becomes another delegation point in the DNS infrastructure tree. The correct operation of the delegation hierarchy is essential to the stability of the DNS.

There are two fundamental requirements for the correct operation of the delegation. First, the parent of a sub-zone must point to the sub-zone server(s). Second, the sub-zone server(s) have to be in operation and be reliable.

Extrapolating from various DNS specifications "correct operation" can, for the purposes of this recommendation, be defined as follows:

1. A parent zone is responsible for defining its points of delegation (sub-zones), ensuring the availability of a server to respond with the appropriate referrals (NS records) for those sub-zones, updating the referral information upon request from the child sub-zone in a timely fashion, and ensuring the child honors its responsibilities.
2. The child zone is responsible for ensuring the availability of a server to respond to queries about it and ensuring its parent zone is up-to-date with respect to the referral information it maintains on behalf of the child.

The first requirement is straightforward and unambiguous. If there is incorrect information in a zone's parent, the zone is unreachable and effectively out of service. Therefore, the zone's parent must have the correct referral information for the zone. Whenever a zone server is relocated to another address, the zone's parent must be updated forthwith.

The second requirement is less crisp. Consistent with past recommendations, we recommend that each zone operate at least two independent servers to provide a high degree of reliability and availability. This recommendation applies most strongly to the root and top-level domains, but we also recommend that zones with a high volume of DNS queries and/or zones that aspire to be highly available also operate two or more independent servers.

It is not always easy to determine if a set of servers is mutually independent. In addition to using different hardware, the servers ideally should be on different networks and in different physical locations.

1 Introduction

The DNS infrastructure can be defined by its points of delegation. Beginning with the root of the DNS ("."), each zone administrator has the authority to delegate sub-zones to other responsible parties. Each sub-zone becomes another delegation point in the DNS infrastructure tree. The correct operation of the delegation hierarchy is essential to the stability of the DNS. Applications routinely query about unknown-to-them domain names by querying first at the root, recursively querying each of the points of delegation according to the referrals it receives, and continuing until it reaches the server of the zone with the desired information.

One of the requirements necessary to ensure the correct operation of the delegation hierarchy is to ensure that a zone's parent has the correct referral information with which to respond to queries about the zone. A zone and its parent must work together to ensure the parent always has the correct referral information and the parent must update the referral information upon request in a timely fashion. If a zone's parent responds to queries with incorrect referral information the zone's servers will be unavailable just as if they were not present at all.

A second requirement necessary to ensure the correct operation of the delegation hierarchy is to ensure the availability of a server at each of the points of delegation (at least one server for each zone). When there is no server available for a zone then all applications and services that depend on the DNS to locate other Internet services in that zone and all that zone's sub-zones can not complete their task. The DNS was originally defined to require at least two independent servers for each zone with mechanisms to synchronize data between those servers for precisely this reason.

Extrapolating from the DNS specifications (RFCs 1034, 1035, and subsequent modifications) and various DNS operations documents (RFC 1591 in particular), "correct operation" can, for the purposes of this recommendation, be defined as follows:

1. A parent zone is responsible for defining its points of delegation (sub-zones), ensuring the availability of a server to respond with the appropriate referrals (NS records) for those sub-zones, updating the referral information upon request from the child sub-zone in a timely fashion, and ensuring the child honors its responsibilities.
2. The child zone is responsible for ensuring the availability of a server to respond to queries about it and ensuring its parent zone is up-to-date with respect to the referral information it maintains on behalf of the child.

2 Discussion

The first requirement to ensure the correct operation of the delegation hierarchy is the maintenance of the referral information at the point of delegation. A parent and a child need to cooperate to ensure that the parent has the correct referral information with which to respond to queries about the child zone.

The stability of the DNS infrastructure is directly affected by the failure of either a parent or a child to maintain the correct referral information. When a child fails to update its parent when its referral information changes the child risks being disconnected from the Internet. Since the child is independent of the parent and the child is solely responsible for its zone, the parent can not be held responsible if the child finds itself disconnected from the Internet because the child failed to inform its parent of a change in its referral information. A parent may, if appropriate and necessary based on criteria outside the scope of this recommendation, choose to move the delegation of a particular sub-zone in order to improve its availability.

Conversely however, a child may not have the option of moving its zone if a parent fails or chooses not to fulfill its responsibility to update the referral information upon request in a timely fashion. In many cases a child zone may be obligated or restricted to one particular parent. The stability of the DNS infrastructure requires that the parent update the referral information upon the request of the child. Failure to do so will result in the child being disconnected from the Internet. It follows that unless the objective of the parent is to disconnect the child from the Internet, for reasons that are beyond the scope of this recommendation, the parent should always fulfill its responsibility to update the referral information. In general, a child and parent should have a clear and unambiguous understanding of the criteria with which a parent will refuse to update a child's referral information.

With respect to the requirement for the availability of a server for a zone, this recommendation focuses on two issues: whether or not two independent servers are always required and precise definition of "independent." For example, consider the case of a zone that does not delegate any sub-zones.

When a zone has no sub-zones then the only party affected by a failure of the zone's server is the zone itself. Without sub-zones the DNS delegation hierarchy is unaffected by the single point of failure represented by the single server serving that zone.

On the other hand, a zone that supports the delegation of sub-zones to distinct responsible parties has a greater responsibility to ensure there is a server available since the number of affected parties is directly proportional to the number of sub-zones. The requirement for multiple independent servers is applicable in this scenario.

Specifically, in this context "independent" is defined as follows. An operational policy that increases DNS infrastructure stability for zones that delegate sub-zones would be one in which at least two servers are present for the zone and for which the following is true:

1. The servers are physically located in geographically different sites.
2. The servers are physically connected to the Internet through different paths, which means either their upstream network provider is different or their respective paths to the upstream network provider are physically different.

More succinctly, a parent zone should ensure that the servers for both itself and each of its child zones that also have delegated child zones are not behind a single point of failure. Such points include the following.

- The network infrastructure, e.g., a single LAN, a single router, a single service provider, etc.
- The physical infrastructure, e.g., a single co-location facility, a single geographical location, etc.

3 Confirming Multiple Independent Servers

Unfortunately, there is no completely automated method for determining if a zone employs multiple independent servers. There are a number of automated checks that can be employed to provide a partial answer to the question.

1. Are there at least two NS records for the zone?

   Although a parent zone will have a copy of this information for each sub-zone that it delegates, it is useful to ensure the correctness of the information by querying the servers for the sub-zone directly and comparing the information with that of the parent.

   If the sub-zone does not have two NS records then it is known that independence is not present. Although a zone must have two NS records in order to have independence, the presence of two NS records is inconclusive.

2. Are the IP addresses of at least two NS servers on different networks?

   Given the NS records it is straightforward to query the DNS again to obtain their IP addresses and observe if they are on the same network or not.

   If the IP addresses are on the same network then it is known that independence is not present. Although the IP addresses must be on different networks to have independence, an IP address only provides a virtual reference to its physical network topology and therefore the presence of different networks is inconclusive.

3. Are at least two IP addresses in different Autonomous Systems (ASs)?

   If the IP addresses are in different ASs then network connectivity independence is known to be present. However, no conclusion can be stated regarding geographic independence. If the IP addresses are not in different ASs then no conclusion can be stated.

There is no automated method to determine if geographic independence is present. Checks 1 and 2 must be true but in addition it is necessary to manually confirm the physical location of the servers in question.

If Check 3 is false then the only way to determine if network connectivity independence is present is to manually confirm the physical configuration of the servers in question and the network to which each server is physically connected.

4 Summary And Recommendations

This recommendation only addresses the stability of the DNS infrastructure with respect to two issues: the timely update of referral information in a zone's parent and the requirement for two independent servers for every zone.

Specifically, because the Root zone (".") and the Top Level Domain (TLD) zones both delegate sub-zones to multiple independent parties, the correct operation of their servers is critical to the stability of the DNS infrastructure. As described by the discussion above, "correct operation" suggests the following two recommendations.

1. Zone administrators should adopt a policy that ensures that referral information for their sub-zones is updated upon request and in a timely fashion.

2. Zone administrators should adopt a policy that requires multiple independent servers for their zone when it delegates sub-zones to more than one responsible party.

It is worth noting that neither a zone transfer nor access to a zone file is required or sufficient to verify the stability of the DNS delegation hierarchy.

5 Acknowledgements

The members of the Security and Stability Advisory Committee prepared this document.

Alain Aina (Consultant)
Jaap Akkerhuis (SIDN)
Doug Barton (Yahoo!)
Steve Bellovin (ATT)
Rob Blokzijl (RIPE)
David Conrad (Nominum)
Steve Crocker (Shinkuro), Chair
Johan Ihrén (Autonomica)
Mark Kosters (VeriSign)
Allison Mankin (Lucent Bell Labs)
Ram Mohan (Afilias)
Russ Mundy (SPARTA, Inc.)
Jun Murai (Keio University)
Frederico Neves (registro.br)
Ray Plzak (ARIN)
Doron Shikmoni (ForeScout)
Ken Silva (VeriSign)
Bruce Tonkin (Melbourne IT)
Paul Vixie (ISC)
Rick Wesson (Alice's Registry)

Support for the committee is provided by Jim Galvin (eList eXpress).