Upper limit values for DNS

draft-fujiwara-dnsop-dns-upper-limit-values-01

Kazunori Fujiwara, JPRS

IETF 121, dnsop WG

Upper limit values: Problem Statement

- Some parameters in DNS don't have clear upper limits
  - Number of Resource Records in an RRset
    - Number of RRSIG/DNSKEY/DS RRs in an RRSet
    - Number of NS, glue, …
  - Number of CNAME/DNAME chains
  - Number of levels of unrelated only delegations
  - DNS packet size (≦65535)
- Without upper limits,
  - Easy to attack DNS aimed at resource depletion or DoS
    - Just prepare long CNAME chains, large RRSets (many RRs) in a zone
  - Several attack methods have been reported
    - KeyTrap, Tsuname, several DoS attacks
- This draft proposes reasonable upper limits for DNS protocols
  - Intended status is "Best Current Practice" (or update DNS standards ?)

Possible upper limit items

Name		use cases	implementation
DNS message size		< 65536	
Number of Resource Records in a RRSet	13 ?	. / com NS	
Number of NS RRs at a delegation	13 ?	. / com NS	
Number of glue RRs at a delegation	26 ?	com glue	
Number of DS RRs at a delegation		need research	
Number of DNSKEY RRs in a DNSKEY RRSet		need research	
Number of RRSIG RRs for each name and type		need research	
Number of levels of unrelated only delegations		need research	
Number of CNAME/DNAME chains		10	

Details: Packet size limits

- There were comments that there are size limitations even if no precise upper limit is set.
- The DNS packet format has an upper limit of 65535 octets, so an RRset cannot exceed that size
  → Attackers use this size to carry out resource-wasting attacks
- The size of a DNS response that can be sent using unfragmented UDP is about 1400 octets (see avoid-fragmentation draft)
- This size 1400 may be a good limit of an RRSet + RRSIGs + DNS header + question section (if we use UDP)

Details: Number of Resource Records in a RRSet

- Best Current Practice documents should allow for values that are currently in widespread use.
- Since there are 13 root name servers and 13 name servers for com and net TLDs, the maximum number of NS RR in an NS RRSet should be larger than or equal to 13.
- BIND 9 introduced 'max-records-per-type' parameter and the default is 100.
  - CVE-2024-1737 "BIND's database will be slow if a very large number of RRs exist at the same name
- My proposal: Number of Resource Records in a RRSet ≦ 13

Details: DNSKEY, DS, RRSIG

- KeyTrap is a vulnerability caused by the fact that there is no upper limit on the number of DNSKEY, DS, or RRSIG RRs
- Considering the DNSKEY rollover and the multi-signer model, the maximum number of DNSKEYs may be 6.
  - (signer 1's KSK, ZSK, new KSK, new ZSK + signer 2's KSK + ZSK)
- Maximum number of DS is 3 (3 KSKs)
  - Multiple Digest Types ?
- Unbound introduced the maximum number of RRSIG validations for an RRset (MAX_VALIDATE_RRSIGS) as 8
  - My preference is 2 (DNSKEY RRSIG may be 3 ?)

Details: Number of alias levels using CNAME/DNAME

- Many resolver implementations can resolve over 10 CNAME aliases
- Unbound introduced 'max-query-restarts' parameter and the default is 11
  - Hard limit on the number of times Unbound is allowed to restart a query upon encountering a CNAME record
- my preference is 1
- However, many domain names use 3 CNAME chains

Details: unrelated only delegation levels

- Unrelated name server names are required for DNS hosting services.
  - for example, ietf.org's name servers are {ken,jill}.ns.cloudflare.com
- However, using unrelated name server names increases the name resolution costs and may increase the likelihood of name resolution errors.
- To avoid complex name resolution and misconfigurations, it is better to avoid using unrelated name server names as much as possible
- The draft proposes to use in-domain name servers as much as possible for name resolution of unrelated name server names
- Unrelated name server names SHOULD be hosted by a domain name with at least one in-domain name server name.
- In other words, DNS provider SHOULD have at least one in-domain nameserver for their domain names.
  - cloudflare.com's name servers are ns*.cloudflare.com (in-domain only, excellent)

Proposed upper limits

Name	proposal	use cases	implementation
DNS message size (without PQC)	≦ 1400		≦ 1232 on UDP
Number of Resource Records in a RRSet	≦ 13	. / com NS	≦100 (BIND)
Number of NS RRs at a delegation	≦ 13	. / com NS	
Number of glue RRs at a delegation	≦ 26	com glue	
Number of DS RRs at a delegation	≦ 3 ?	need research	
Number of DNSKEY RRs in a DNSKEY RRSet	≦ 6 ?	need research	
Number of RRSIG RRs for each name and type	≦ 2 ?	need research	≦8 (Unbound)
Number of levels of unrelated only delegations	≦ 2 (1)	need research	
Number of CNAME/DNAME chains	≦ 3	10	≦11 (Unbound)

Recursive resolvers SHOULD respond with a name resolution error (Server Failure) if they receive responses from authoritative servers that exceed these limits.

Request

- This draft proposes aggressive upper limits to advance discussions on determining upper limit values in DNS protocol.

- Please review draft-fujiwara-dnsop-dns-upper-limit-values