Council of European National Top Level Domain Registries

DNS Easy 2011

Rome, Italy

October, 18th 2011

I. Less secure at DNS security Conference

DNSSEC and current DNSEC problems, the ISC reputation system and models and tools to measure the health state of the DNS were discussed during the DNS Easy 2011 Conference in Rome this week. The Conference, jointly organized by ICANN, DNS OARC and the Global Cyber Security Center (GCSC) as the local host, preceded the 3rd Global DNS Security, Stability & Resiliency Symposium after meetings in Atlanta and Kyoto. The symposiums were watched with some concerns by some registry operators who feared that ICANN would be interested in a bigger role in DNS security. But talks about a DNS CERT, for example, seem to have died down in the meantime and were not raised during the DNS EASY day.

The DNS EASY day was an open workshop to present new pieces of research and some high-level key-note speakers like DNS pioneer Paul Mockapetris (Nominum) and Paul Vixie (ISC). The SSR Symposium was on invitation only and obliged participants to the Chatham house rules, not allowing person-attribution of what was said during the event.

The GCSC presented by its Director General, Andrea Rigoni during the welcome session, is a pretty new foundation dating from May, 7th, 2010, with Poste Italiane being the founder (founding capital 2,4 Million Euro, see the first balance report for 2010 here). According to Rigoni Enel Group, Master Card and Almaviva are also participating members in GCSC. One of the partners listed is the Department of Homeland Security. Poste Italiane's CEO, Massimo Sarmi, is Chairman of the GCSC Board.
Poste Italiane, according to media reports, has been one of the most active players in Europe's anti cybercrime efforts and also initiated the setup of the European Electronic Crime Task Force, which in effect is a bi-lateral initiative by Italy (Italian post and Italian law enforcement and the US Secret Service).

Given this background there were many ironical remarks by experienced DNS experts at the first day of the conference about the conference network, which could only be accessed when disabling DNSSEC (one participant said deep packet inspection was also used to control the traffic) and needed hacking around blocked ports to allow the use of SSH or VPN clients. Participants from the registries joked about the paradox that Italian security laws made the net in fact less secure. According to Rigoni the network in fact was strictly set-up according to Italian laws. Every participant got a dedicated password.

II. DNSSEC failure, figures and counter-strategies

Several papers looking into DNSSEC failure rates were presented during DNS Easy day. The Japanese registry JPRS and the National Institute of Informatics in Japan tried to quantify the impact of DNSSEC validation failure (through key expiration, key roll-over, key registration – the paper listed five well-known events at RIPE NCC, Nominet, Mozilla and IAB in 2010 and AFNIC in 2011) on cache servers and autonomous systems (AS).

Kensuke Fukada from the NII presented the results of a 24 hour measurement of how fast and how far a DNSSEC validation failure of an authoritative DNSSEC server would be propagated (data analysed was based on tcpdump of DNS traffic to and from seven .jp servers and their instances, altogether 1.4 billion queries, but with no real DNSSEC traffic, as registration of DS keys was only offered at a later time).

The study found that 18 percent of cache servers would fail to validate DNSSEC in the first 10 minutes if a single failure did occur in a .jp DNS server, 15000 ASes would be affected during the same time. After two hours 85 percent of ASes. 35 percent of IP addresses and 25 percent of all cache servers would fail to validate. Given the current number of validation is still low, the impact at this time would only be visible at 0.8 percent of the total number of cache servers and 12 percent of ASes. After one day 34 percent of the addresses and 75 percent of the ASes would be affected, a quite significant rate, according to Fukada.

The study also checked how fast a wrongly registered DS key was noticed in the net, concluding that a 10 minute failure of a top name resulted in more than a 1000 cache servers failing to resolve the name. Within an hour this number would increase to 4000- 40000 cache servers. Early detection could be realized by either using cache servers as sensors or by monitoring the validity of one's TLD with actives probes at their cache servers.

Mitigation on the other hand was much more difficult – recovery times according to the study ranged from hours to days. Removing or adding DS records to the parent according to the study was 15-30 minutes for .jp, but 2-3 days for the root. The latter was a rather optimistic figure, Paul Vixie said to this reporter. As there had to be paper exchanges between ICANN, the Department of Commerce and VeriSign the time could be considerably longer. Paul Mockapetris' reaction was that possibly a “stand-by” key had to be in place for quicker mitigation at the root zone. During his keynote talk Mockapetris also noted that the experts seemed to be all satisfied with the centralized trust anchor design and thought that multiple roots might be inconvenient. But himself, he considered multiple trust anchors as fundamental. Mockapetris nevertheless said he hoped for DNSSEC to get “a few miles on the tire” and be widely implemented by 201x (first decade).

A rather pessimistic view on DNSSEC failures today and in the near future was presented by Casy Deccio, Senior Member of Technical Staff and Sandia National Laboratories, author of DNSViz (a web based tool for DNS analysis). Sandia National Lab is a government-owned/contractor operated (GOCO) facility managed by Lockheed Martin for the US Department of Energy's national Nuclear Security Administration. As such, it is in the .gov-zone, a zone that saw a high number of DNSSEC validation failures. Deccio said that a lot of problems resulted from “people getting in to DNSSEC and then letting it go” as they did not have time to focus on the risen demands on maintenance. Deccio in his paper listed six common types of misconfiguration:

> ***DS mismatch***, *DS RR are present in a parent zone, but none correspond to any self-signing DNSKEYs in the child zone; (382 events)*
>
> ***DNSKEY missing***, *DNSKEY is referenced in the RRSIG or DS, but not included in the DNSKEY RRset (1250 events, 1 TLD affected)*
>
> ***NSEC missing***, *lack of NSEC RRs in a negative response (non-existent domain name, 449 events, 3 TLDs affected)*
>
> ***RRSIG missing***, *authoritative server does not provided the RRSIG (2418 events, 6 TLDs affected)*
>
> ***RRSIG bogus***, *signature in record data of RRSIG does not validate against the RRset it covers (284 events)*
>
> ***RRSIG dates***, *if RRSIG is allowed to expire, or is published before its inception date (1691 events, 4 TLDs affected)*

What made Deccio's presentation look so pessimistic were the failure rates found in a survey of 2242 production signed zones between June 2010 and July 2011. “For signed zones under the .gov TLD, which made up the largest contingency of those analysed, over 40 percent of the zones experienced some type of misconfiguration. For nearly all TLDs shown, at least 30 percent experienced some type of misconfiguration.” The high rate possibly should be qualified as experts by several registries present pointed for example to incomplete data with regard to the signed zones. For .cz top level domain, for example, the number of signed SLDs was much higher than captured in the survey, and while there had been a considerable number of problems (about 1000, due to registrar changes), given the actual number of signed zones (according to the counter on the CZ.NIC-site at the time of this writing: 144852) the percentage should be much smaller.

Nevertheless there have been discussions between registry experts about how to eliminate some sources for failure, roll-over of non-compromised keys just as a regular exercise should be re-examined, for example.

A long list of automation tools and monitoring tools (the DNSSEC-Tools suite) to facilitate the necessary DNSSEC maintenance was presented by Robert Story from Sparta. Operations that could be automated according to Story are: zone re-signing and key management, key generation, algorithm roll-over, key roll-over, DS transfers to parent, trust anchor roll-over.

III. Challenges to the DNSSEC development

Rick Lamb, responsible for DNSSEC at ICANN, in his talk called for a race to the top in DNSSEC deployment and portrayed DNSSEC as the basis for a long list of new options delivered by a secure DNS infrastructure, from self-certificates (developed in the DANE WG a the IETF), to secure email (S/MIME) to identify services, that were desperately looked for by governments. While all these are options, the uptake of DNSSEC so far has been very slow, with only 1 percent of domains signed out of 81 percent for which DNSSEC was available (72 TLDs out of 310, including all large TLDs, are signed by now).

Lamb has one concern bigger than the slow uptake and that is launch failure resulting from turning DNSSEC on without additional resources. If registrars offered DNSSEC to sign one's domain, the ISP just turned validation on, but then goes back to the classical DNS mentality of “set and forget”, this may lead to a corrupt DNSSEC service, downtimes resulting from expired signatures and inaccessibility from a validating ISP's network. This in turn could very well result in ISP stopping validation again and DNSSEC being “frozen in a deep hole” from where recovery would take years.

IV. Reputation systems, not everybody's darling

Bind-guru Paul Vixie (ISC) announced the the 3.0 version of “Response Policy Zone (RPZ)” that is already available in BIND 9.8.X and allows the (live) blocking of malicious/dangerous/unwanted source domains at the level of a recursive server based on lists drawn from blacklist providers and on policies about how to deal with these domains (return a fake alias, for walled gardens, return a fake NXDOMAIN, to blackout the name, return a fake answer to the type being queried, protect the name against subsequent policy triggers). Vixie has started to promote RPZ since 2010 when he posted an article about “taking back the DNS” or using the recursive DNS server as a “security hammer”.

All abusers had to have some assets like botnets, domain names IP addresses, name servers and all could serve as a trigger to “refuse to talk to them”, Vixie explained. The core concept is shown in the graph below (find complete ISC presentation here).

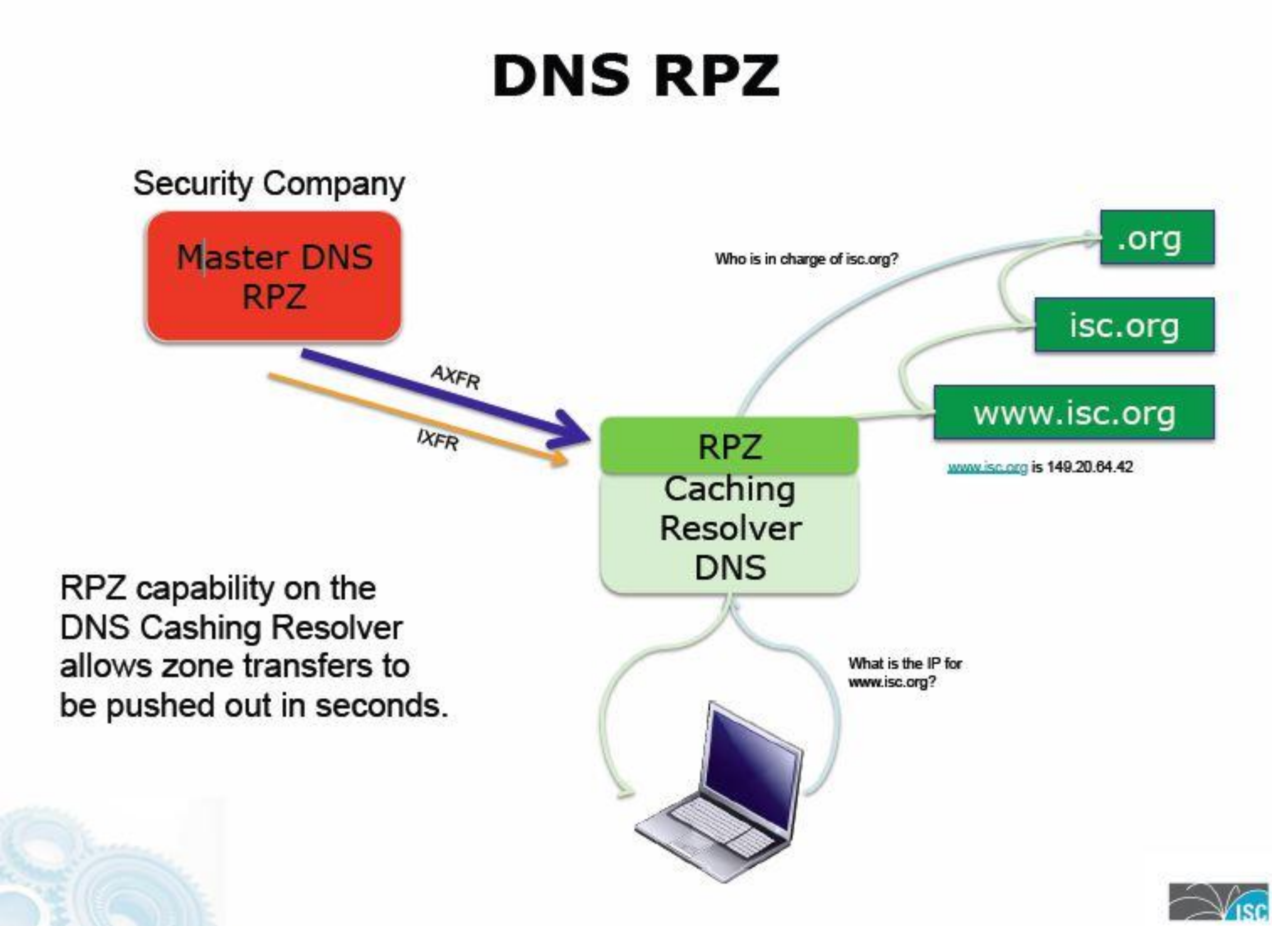

DNS RPZ

RPZ capability on the DNS Cashing Resolver allows zone transfers to be pushed out in seconds.

Security Company
Master DNS RPZ
AXFR
IXFR
Who is in charge of isc.org?
.org
isc.org
www.isc.org
www.isc.org is 149.20.64.42
RPZ
Caching Resolver DNS
What is the IP for www.isc.org?
ISC

Several blacklist providers (Spamhouse and SBRL) already have or are working on RPZ-compatible lists. Yet he predicted that “DNS firewalling will become ubiquitous” in the same way that anti-spam technologies became ubiquitous. RPZ would “create a global market” with hopefully 400 instead of 40 “policy providers”, said Vixie, and interfaces available from all DNS providers, not only from ISC. He said it was the only way to keep keep up with the speed of the attackers.

While Vixie was very optimistic about the effects of the RPZ use, in earlier presentations he himself has pointed to potential problems, for example a potential interest from governments to use the RPZ for censoring content. Some RPZ data sources might be politically, racially, or religiously motivated, Vixie warned. From an architectural view Vixie also has acknowledged that “as with all reputation systems, the systemic effect on the DNS will be to make it less reliable”. In Rome, Vixie pointed to open questions about how to make sure that DNSSEC validated malicious sites could still be put on the black list. In general the “interaction” between DNSSEC and RPZ still had to be worked on.

Representatives from several registries confirmed that there were concerns about what can be seen as a dual use-nature of such tools.

The “filtering” topic (and its relation to DNSSEC) was also discussed on day one of the SSR Symposium. When an infrastructure for filtering was provided, the “bad guys”, too could use it more and more. With regard to the potential erratic effects (or lack of effectiveness) experts seem somewhat cynical about many governments satisfied with “security theater”. In Italy the local prosecutor for example faxes a list of about 500 domains that have to be blocked, and the blocking mainly has an effect on the less savvy, ordinary users that do not use alternative DNS providers for example.

DNSSEC was no problem for censors as long as they just want to make content inaccessibly, but it makes the filtering obvious. A last trick for “liars” could be to drop the requests (something which will be done for example by China Telecom against VPN servers, see below). An optimistic scenario, that filtering would stop when DNSSEC was widely implemented, is highly unlikely according to the experts. A rather negative effect will be that stub validation could become impossible and DNSSEC then will be limited to caches of ISP. This will make it difficult or impossible for new DNSSEC based applications to be developed.

Mockapetris told the conference in his talk that he was convinced that DNSSEC and reputation systems (as part of simpler security measures) had to be brought to the edge, to the end user. Engineers had to get the interfaces much more user-friendly and really simple.

V. Model for predictions on DNS

There is a lot of interest to better and better monitor DNS health and possibly predict the impact of new things arriving like DNSSEC, new gTLDs or more filtering. Bart Gijsen from the Dutch Research Institute TNO presented a “global reference model for the DNS” that shall allow to make predictions about the impact of new developments. The model that was based on a classical client, resolver, authoritative server-architecture and was developed from figures of the Dutch provider SurfNet, experiments with DNS servers and data from sources like “a day in the life of the Internet” (CAIDA). They were validated against data from another source.

According to Gijsen the reference model will be refined further, but it already showed promising rates of accuracy with regard to predictions. The test taken for the prediction of the impact of a new development was a scenario with increased servfail answers due to second level domain DNSSEC configuration error.

Gijsen reported that the Dutch representative in the Governmental Advisory committee of ICANN was very interested in using the model. Governments have over the recent month always asked for predictions about what will happen when several hundred new TLDs are introduced to the system.

VI. Take the DNS's pulse

A new project to measure DNS health (Mensa) of the DNSEASY host, the Global Cyber Security Center (GCSC), was presented in Rome by João Damas. The target for Mensa, according to Damas' talk, is aggregation of different DNS views (from different places, organisations). The result would be an index that would give a rough health evaluation. Works seems to be only starting: Damas said that a test-bed now has been set up, but more aggregation was necessary.

Representatives from CZ.NIC presented work on a DNS traffic anomaly detector which allows to detect both low-volume anomalies (“a distinct pattern of an almost unique set of a few hundred resolver IP addresses concentrated into a few seconds” for a foreign embassy's A resolver) and high-traffic anomalies (scanning). “Packet attributes which are selected to serve as hash keys affect the classes of anomalies which are identified by the tool. Two policies implemented with different effect: a) IP policy serves best for domain enumeration detection b) Query name policy divulges domain-related events (e.g. presence of short TTL domains, like in fast flux).” The tool is open source and available from CZ.NIC website. It shall allow to pinpoint suspicious abnormal or just interesting traffic in real-time, but analysis has to be done manually.

VII. China Telecom trying to suppress “black” VPNs

A presentation judged as “alarmingly good” by participants came from China. Liu Zhiqian, according to his bio the leading architect for China Telecom's WIFI network, reported about an analysis of “free riders” using the Loopc VPN software to access the Internet over China Telecom's network. Loopc can be bought online on Taobao shopping portal for a monthly price of 28 Yuan per month (prices for one day were also possible when this reporter visited the Taobao site first). Liu said his lab had bought a version for the test and then analysed what he called the “tricks” used by the software to deliver free Internet access.

The Chinese engineer described three different variants of the IP over DNS set-up offered by Loopc. In all cases China Telecom's broadband remote access servers (BRAS) are used to connect to a Loopc proxy. Version one allowed connection from the BRAS directly to one of the 12 Loopc servers exploiting a vulnerability, namely the possibility to send packets directly to the BRAS before authenticating over a captive server.

If the first DNS request is not forced to go to the provider's recursive name server, the connection to the Loopc proxy can be established and Loopc server and user will exchange what are seemingly malformed DNS messages over the BRAS server. Another variant observed by the China Telecom experts is the tricking of the providers DNS server to pass on the DNS requests to the Loopc server in case the BRAS server forces to send DNS requests that way.

> *“The client encodes the user's HTTP requests by the RSA key, encapsulates them into DNS queries for names in cry555.com (Domain for Loopc server) and sends them to the provider's name server. After receiving these queries and looking through the cache, the providers names server finds no hit and starts iterative queries. These queries are then received by the authoritative server for 555cry.com which is actually the Loopc VPN proxy server. The Loopc VPN server will decode the queries and fetches for the client the desired contents. After that the Loopc VPN server encodes those contents and encapsulates them into pseudo DNS responses which are then transmitted to the client through the BRAS.”*

Still under development, according to Liu, is variant three, which involves the provider's recursive name server in both directions. The software Iodine did this already, and Loopc seemed to be developed in this way, Liu said. To prepare against this development China Telecom proposes to use blocking based on “Suffix Matching/Requested Frequency Counting” (SM/RFC as bi-directional access control lists on the name server are of no help).

The SM/RFC algorithm counts the number of queries for a suffix and if a the count exceeds a threshold a penalty value is assigned, and further if a suffix penalty value accumulates to a certain value, recursive queries to the domain will be dropped. China Telecom whitelists big commercial providers like Baidu to prevent queries to them to be dropped. Whether one can apply to be put on this whitelist is an open question. Liu said the system was currently being tested.

While Liu said the major goal was to stop users routing around billing while eating up capacity in the network – and this being a commercial venture – the main motivation for users in China still might be not so much a commercial benefit, but might stem from the other connotation of “free access”.