OARC 28 (San Juan)

Report of Contributions

Contribution ID : **40** Type : **Standard Presentation**

Measuring Efficiency of Aggressive Use of DNSSEC-Validated Cache (RFC 8198)

Thursday, 8 March 2018 14:30 (30)

In this presentation we will analyze data from real recursors to quantify impact of RFC 8198 on real traffic. Was is worth the effort, or is it waste of energy to implement it?

Summary

Talk Duration

30 Minutes

Primary author(s) : Mr. ŠPAČEK, Petr (CZ.NIC)

Presenter(s) : Mr. ŠPAČEK, Petr (CZ.NIC)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **41** Type : **Standard Presentation**

DNS infrastructure at Facebook

Thursday, 8 March 2018 11:30 (30)

At Facebook, we leverage the DNS for multiple purpose, internally, we use it to access servers via hostnames, service discovery and load balancing. Externally, our authoritative nameservers are helping us in steering the traffic of the people using our products to a point of presence that will provide them with the best experience.
With constant churn in our infrastructure, cluster coming up and down, machines being re-provisioned, traffic patterns, maintenances, capacity and potential outages, keeping DNS mapping up to date, accurate and with minimal to no operator intervention is challenging, so it is to ensure consistent and fast distribution of those maps across our fleet of DNS servers.
This talk will explain how we have set up a pipeline that, by leveraging multiple simple and self contained components, open-source software and some python glue, keeps our DNS mapping up to date across all our nameservers.

Summary

Talk Duration

30 Minutes

Primary author(s) : Mr. BRETELLE, Manu (Facebook)

Presenter(s) : Mr. BRETELLE, Manu (Facebook)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **43** Type : **Standard Presentation**

Identifying DNS Open Resolvers in IPv6

Friday, 9 March 2018 11:30 (15)

Introduction

As all of you know, having DNS servers considered Open Resolvers is very negative, both for those who leave the service open, for the Internet and for online security. To read about Open Resolvers I recommend reading this link: https://www.certsi.es/blog/dns

Identifying a DNS Open Resolvers in IPv6 (open DNS servers)

Identifying Open Resolvers servers or open DNS servers in the world of IPv4 is easy, due to the short length of the IPv4 space (2 ** 32) it is relatively easy to run check every IP.
In the world of IPv6 it is virtually impossible to verify each IP address, I mean, to test IP by IP. If we try this test can last thousands of years

How is a DNS Open Resolver identified?

A recursive DNS server should only answer queries to its own clients (yes, there are few exceptions) and should reject any other. For example, the DNS servers of the ACME ISP should only respond to queries from their own clients, to no one else.
Our test consist in querying a domain name (such as www.lacnic.net) to a list of DNS servers, if the DNS server responds with a response then it is considered Open Resolver, if it returns a rejection (Query refused) or simply timed out it is not an Open Resolver.

How we find the list IPv6 resolvers?

Lacnic manages a server that can be called: Reverse Root Server, specifically the letter "D", that is, d.ip6-servers.arpa. Many queries looking information for reverse DNS goes throughout this server, in general this server ONLY receives queries from DNS servers. This is where they get IPv6 addresses from DNS queries. Since this server does not allow recursion every IP that queries this device can be considered a resolver.

Procedure & Algorithm:

.- In the server, we captured over 2.5 millions packets. We captured only v6 packets, filtered by port 53 and destined only to the IP address of the server.

.- Malformed packets, errors, etc. were discarded.

.- From the IPv6 addresses obtained in step 2 a list of unicast IP addresses is created (that is, duplicates are deleted)

.- Finally we got a list of over 800.000 resolvers (not Open resolvers yet)
A python script takes each IPv6 address from the list gotten in item 4, and queries the FQDN www.lacnic.net, verifies recursion and the status of the response. In the case of an Open Resolver, the IP is registered

.- Some manual verification is also performed, I mean, we take some results (text file) and manually check the results.

What is going to be shown

In case this paper is accepted for presentation we are going to show some this mechanism accompanied with some results & statistics

Similar publications

https://labs.ripe.net/Members/luuk_hendriks/finding-open-dns-resolvers-on-ipv6

Partial raw data:

http://stats.labs.lacnic.net/BORRAR-v6-resolvers.txt

Summary

In this presentation we will show a mechanism we used to identify IPv6 Open Resolvers in Internet. We will also present some statistics regarding this fact and possible some things TO DO for the future

Talk Duration

15 Minutes

Primary author(s) : Mr. ACOSTA, Alejandro (LACNIC)

Presenter(s) : Mr. ACOSTA, Alejandro (LACNIC)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **44** Type : **Standard Presentation**

BIND 9 Past, Present, and Future

Thursday, 8 March 2018 16:30 (30)

BIND 9 is now 17 years old, the latest stable version 9.12 was releases in December and the BIND 9 Team has adopted changes to adapt to the ever change Internet landscape.

Summary

In this talk, I will present BIND 9 colorful past, the current state of development, and the changes that BIND 9 Team has adopted to cope with modern DNS. I will talk about the changes in the development model, release cycles, and also about the planned features that will help BIND 9 Team to be more nimble in adding new features, fixing old issues, spending less time on maintenance, while ensuring the stability for existing users. I believe that existing BIND 9 users will be thrilled.

Talk Duration

30 Minutes

Primary author(s) : Mr. SURY, Ondrej (Internet Systems Consortium)

Presenter(s) : Mr. SURY, Ondrej (Internet Systems Consortium)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **46** Type : **Standard Presentation**

Evaluation and consideration of multiple responses

Friday, 9 March 2018 14:30 (15)

Three drafts proposed authoritative servers to respond additional resource records to pre-populate resolvers' cache. The author made an authoritative server patch to add additional A/AAAA/NSEC RR in additional section when the server receives A or AAAA query. This talk reports evaluation result of multiple responses and considerations. It contains brief introduction of multiple response proposals, experimental authoritative server implementation, experiment setup, experiment result of BIND 9, Unbound, Knot Resolver. It also discusses the effectiveness of multiple response drafts.

Summary

Talk Duration

15 Minutes

Primary author(s) : Mr. FUJIWARA, Kazunori (Japan Registry Services Co., Ltd)

Presenter(s) : Mr. FUJIWARA, Kazunori (Japan Registry Services Co., Ltd)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **47** Type : **Standard Presentation**

Using DITL data to look at leaked queries

Friday, 9 March 2018 11:45 (30)

The DITL data collected at DNS-OARC can be used for a variety of research. Here, I analyze QNAMEs in queries to the roots during the DITL 2017 to look at the prevalence of collisions for strings from earlier collision studies (such as "corp" and "home") as well as leakage from TLDs that are not expected to be in the root zone at all. This required looking at the entire dataset, collecting just the QNAMEs, sampling for likely leaked TLDs, and then ranking the data to show which TLDs that were not in the root zone were most commonly seen in the dataset.

In order to do this research, DNS-OARC set up a new server and I created new software to efficiently sample the data. The DITL data cannot be moved from DNS-OARC systems, and because looking at the QNAMEs is quite space-intensive, tradeoffs had to be made in the analysis. I describe that software (which has been published) and show how it can be used by other researchers who are using DNS-OARC systems to analyze various DITL data.

As a complement to the DITL-based research, I ran similar tests on L-root data kept at ICANN. In this research, I found some significant leaks that appear much more often during DITL than at other times. I show that L-root data can be used as a reasonable substitute for DITL data for some research, and suggest that data from other root servers might also be used in this fashion.

Summary

The DITL 2017 data shows which names are likely to be the most leaked from enclaves that use their own TLDs. The results are compared to recent L-root data to show that data from a large root operator can be used in a fashion similar to DITL data.

Talk Duration

30 Minutes

Primary author(s) : HOFFMAN, Paul (ICANN)

Presenter(s) : HOFFMAN, Paul (ICANN)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **48** Type : **Standard Presentation**

Field measurement of Anycast DNS root traffic delivery

Friday, 9 March 2018 12:15 (15)

Using queries issued by our test environment targetted at the root servers we analyse DNS traffic observed at the resolver used by the client, the anycast nodes at which the queries arrived and correlate this to the geographical location of the originating client to analyse the effectiveness of the routing system in delivering traffic to the closest Anycast node.
As a by-product we take a casual, but quantitative, look at the various undelegated TLDs in use by the client's environment DNS environment.

Summary

Using queries issued by our test environment targetted at the root servers we analyse DNS traffic observed at the resolver used by the client, the anycast nodes at which the queries arrived and correlate this to the geographical location of the originating client to analyse the effectiveness of the routing system in delivering traffic to the closest Anycast node.
As a by-product we take a casual, but quantitative, look at the various undelegated TLDs in use by the client's environment DNS environment.

Talk Duration

15 Minutes

Primary author(s) : Mr. SILVA DAMAS, Joao Luis (Bond Internet Systems)

Co-author(s) : Mr. HUSTON, Geoff (APNIC)

Presenter(s) : Mr. HUSTON, Geoff (APNIC); Mr. SILVA DAMAS, Joao Luis (Bond Internet Systems)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **49** Type : **Standard Presentation**

Testing Resolver Implementations of RFC 5011 for the Root KSK Roll

Friday, 9 March 2018 10:30 (30)

As part of the assessment of the risk of rolling the root zone's KSK, Verisign commissioned us to performe tests of the implementation of RFC 5011 support in past and present releases of the three open source DNS resolvers Unbound, Bind, and Knot Resolver with regards to the possible sequences of the roll of the root trust anchor. They kindly allowed us to share our findings.

The presentation will first show our methodology—we used CZnic's Deckard to simulate the full time period of a key roll—and the various scenarios we tested, covering both successful key rolls and possible aborts after starting, as well as typical operational occurrences such as installation after the key roll started, resolver restarts during the roll, or non-writeable state directories.

It will then discuss our findings for each of the resolvers and show how their RFC 5011 support developed over the various releases. As a conclusion, we will try to assess what these findings may mean for the success of a root KSK roll.

Summary

Talk Duration

30 Minutes

Primary author(s) : Mr. THESSALONIKEFS, George (Open Netlabs BV); Mr. HOFFMANN, Martin (Open Netlabs BV)

Presenter(s) : Mr. HOFFMANN, Martin (Open Netlabs BV)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **50** Type : **Standard Presentation**

DNSSEC for a Complex Enterprise Network

Thursday, 8 March 2018 14:00 (30)

This talk will give an overview of our planning and efforts so far to deploy DNSSEC for a large enterprise with a complex infrastructure, involving the services of several managed DNS providers. It will start by outlining our specific requirements and design choices (e.g. signing algorithms, authenticated denial mechanisms, signing of dynamically generated records, key rollover schedules, scaling and performance considerations, etc.). Many prominent managed DNS providers have significant limitations in the extent of their DNSSEC support. We will survey DNSSEC capabilities in several of the managed DNS providers, pointing out where they excel, and where they fall short, based on testing we've performed. We will discuss relevant discussions with the vendors and the status of several feature enhancement requests that we've made. A key challenge is the requirement for supporting multiple distinct DNS providers simultaneously, which further complicates the planned implementation, and we will outline several strategies around this. One additional desired goal of this talk is to stimulate a community discussion of what capabilities need to be widely available in DNS providers for successful DNSSEC deployment at many large enterprises.

Summary

Talk Duration

30 Minutes

Primary author(s) : HUQUE, Shumon (Salesforce)

Co-author(s) : ARAS, Pallavi (Salesforce)

Presenter(s) : ARAS, Pallavi (Salesforce); HUQUE, Shumon (Salesforce)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **51** Type : **Standard Presentation**

ODNS: Oblivious DNS

Friday, 9 March 2018 14:45 (15)

It is well known that DNS leaks information that an Internet user may want to keep private, such as the websites she is visiting, user identifiers, MAC addresses, and the subnet in which she is located. This information can be visible to a 3rd party eavesdropping on the communication between a client and a recursive resolver, or even between a recursive resolver and an authoritative server. As this information is sent to each DNS server, DNS operators can also see clients' information.

While there has been some previous work on increasing privacy in DNS infrastructure, such as DNS Query Name Minimization and DNS-Over-TLS, these approaches do not fully solve the problem. Both of these are steps in the right direction, but neither prevent DNS operators from learning information which domains specific users are interested in. Our work is concerned with a powerful adversary that has the capabilities to: 1) eavesdrop on communications between clients and recursive resolvers, and between recursive resolvers and authoritative name servers, 2) request data (via subpoena/warrant) from any number of DNS operators, 3) maliciously access data at any DNS server.

To address this type of attacker, we present Oblivious DNS (ODNS), which is a new design of the DNS ecosystem that allows current DNS servers to remain unchanged and increases privacy for data in motion and at rest. In the ODNS system, both the client is modified with a local resolver, and there is a new authoritative name server for .odns. To prevent an eavesdropper from learning information, the DNS query must be encrypted; the client generates a request for www.foo.com, generates a session key k, encrypts the requested domain, and appends the TLD domain .odns, resulting in {www.foo.com}k.odns. The client forwards this, with the session key encrypted under the .odns authoritative server's public key ({k}PK) in the "Additional Information" record of the DNS query to the recursive resolver, which then forwards it to the authoritative name server for .odns. The authoritative server decrypts the session key with his private key, and then subsequently decrypts the requested domain with the session key. The authoritative server then forwards the DNS request to the appropriate name server, acting as a recursive resolver. While the name servers see incoming DNS requests, they do not know which clients they are coming from; additionally, an eavesdropper cannot connect a client with her corresponding DNS queries.

As this is ongoing work, we have some future work to continue in this direction. We are starting to implement a prototype of ODNS to evaluate its feasibility; additionally, we plan to measure its performance overhead in comparison to current DNS performance. Our prototype will also provide a way to evaluate the design's security and performance.

Summary

Talk Duration

15 Minutes

Primary author(s) : EDMUNDSON, Annie (Princeton University)

Presenter(s) : EDMUNDSON, Annie (Princeton University); SCHMITT, Paul (Princeton University)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **52** Type : **Standard Presentation**

Update on root KSK rollover (or, We're really doing it this time)

Friday, 9 March 2018 09:30 (30)

After the root KSK roll originally intended for 11 October 2017 was postponed because of newly available trust anchor data reported by RFC 8145-capable resolvers, ICANN undertook an investigation to better understand that data and develop a plan for going forward. ICANN has collected community feedback and will publish that plan in early February. This presentation covers ICANN's findings regarding the RFC 8145 resolver data as well as next steps, and offers an important venue for ICANN to collect additional community feedback from DNS-OARC members and workshop attendees.

Summary

Talk Duration

30 Minutes

Primary author(s) : LARSON, Matt (ICANN)

Presenter(s) : LARSON, Matt (ICANN)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **53** Type : **Standard Presentation**

There is always time for DNS

Thursday, 8 March 2018 17:00 (30)

Background

Over the last two years I've been working on reducing the duplicated code that exists between OARC software.

This created a bunch of helper libraries, code that can be added as a git submodule to each software, to handle PCAPs, config files and more.

These libraries also helped creating drool, DNS Replay Tool, and whatever functionality that was missing got added and more helper libraries was created.

Problems

While developing drool I ran into a rather tricky problem, how do you easily configure different scenarios such as read from PCAP A and send to target B but also target C and maybe only for some queries or copies… and so on. In the long run the configure language would more or less become a script language.

I have also not been so happy about the dynamic loadable plugins design in dnscap, they are not easy for anyone to make and a bit of a hassle to package. Thoughts here have been to replace it with some integration to some kind of script engine.

Solution(?)

Many other DNS and packet generating software are using Lua in some form so I took a week to try it out and see if I could put a script engine in drool to tie together all the components.

This gave birth to **dnsjit**!

And I will present more details on how it works and what the future plans.

Summary

dnsjit - Engine for capturing, parsing and replaying DNS

Talk Duration

30 Minutes

Primary author(s) : Mr. LUNDSTRÖM, Jerry (DNS-OARC)

Presenter(s) : Mr. LUNDSTRÖM, Jerry (DNS-OARC)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **54** Type : **Standard Presentation**

The Effect of DNS on Tor's Anonymity

Friday, 9 March 2018 15:30 (30)

Previous attacks that link the sender and receiver of traffic in the Tor network ("correlation attacks") have generally relied on analyzing traffic from TCP connections. The TCP connections of a typical client application, however, are often accompanied by DNS requests and responses. This additional traffic presents more opportunities for correlation attacks. Our work quantifies how DNS traffic can make Tor users more vulnerable to correlation attacks. We investigate how incorporating DNS traffic can make existing correlation attacks more powerful and how DNS lookups can leak information to third parties about anonymous communication. We (i) develop a method to identify the DNS resolvers of Tor exit relays; (ii) develop a new set of correlation attacks (DefecTor attacks) that incorporate DNS traffic to improve precision; (iii) analyze the Internet-scale effects of these new attacks on Tor users; and (iv) develop improved methods to evaluate correlation attacks. First, we find that there exist adversaries that can mount DefecTor attacks: for example, Google's DNS resolver observes almost 40% of all DNS requests exiting the Tor network. We also find that DNS requests often traverse ASes that the corresponding TCP connections do not transit, enabling additional ASes to gain information about Tor users' traffic. We then show that an adversary that can mount a DefecTor attack can often determine the website that a Tor user is visiting with perfect precision, particularly for less popular websites where the set of DNS names associated with that website may be unique to the site. We also use the Tor Path Simulator (TorPS) in combination with traceroute data from vantage points co-located with Tor exit relays to estimate the power of AS-level adversaries that might mount DefecTor attacks in practice.

Summary

Previous attacks that link the sender and receiver of traffic in the Tor network ("correlation attacks") have generally relied on analyzing traffic from TCP connections. The TCP connections of a typical client application, however, are often accompanied by DNS requests. We quantify how DNS traffic can make Tor users more vulnerable to correlation attacks.

Talk Duration

30 Minutes

Primary author(s) : GRESCHBACH, Benjamin (KTH Royal Institute of Technology); ROBERTS, Laura (Princeton University); WINTER, Philipp (Princeton University); PULLS, Tobias (Karlstad University)

Co-author(s) : FEAMSTER, Nick (Princeton University)

Presenter(s) : ROBERTS, Laura (Princeton University)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **55** Type : **Standard Presentation**

Negative Trust Anchors

Thursday, 8 March 2018 15:30 (30)

This presentation will go over some of the issues and basic processes that happen at a large ISP in regards to implementing Negative Trust Anchors.

This will include going over:

- How to determine DNSSEC is broken.
- Determine severity of allowing failed site to stay in state
- Process on when to put in an NTA
- Basic automation efforts around implementation of NTA

Summary

Talk Duration

30 Minutes

Primary author(s) : Mr. CROWE, Joseph (Comcast)

Presenter(s) : Mr. CROWE, Joseph (Comcast)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **57** Type : **Standard Presentation**

The Curious Case of the Crippling DS record

Thursday, 8 March 2018 15:00 (30)

When a DNSSEC Key Signing Key (KSK) is rolled, the Delegation Signer (DS) records in the parent are updated as well. A DS record contains the "Digest Type" used to produce the digest over the KSK. Care must be taken when "rolling" the digest type during a KSK roll. It may well cause the entire zone to become bogus.

My presentation will show how a Top Level Domain went unreachable due to an obscure requirement in the standard and will show inconsistencies between validator implementations.

Summary

Talk Duration

30 Minutes

Primary author(s) : ARENDS, Roy (ICANN)

Presenter(s) : ARENDS, Roy (ICANN)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **58** Type : **not specified**

Disappearing Choice of Recursive DNS Services in Home Networks

Thursday, 8 March 2018 12:00 (30)

Virtually all client devices in homes connected to the Internet obtain recursive DNS server settings automatically. A home LAN can be expected to provide DHCP service, and DHCP can be expected to provide DNS servers that provide some minimal baseline of DNS service.

While many Internet users are completely unaware of the critical service that the DNS provides, a technically savvy user capable of making an informed choice of recursive DNS providers frequently encounters friction actually making that choice due to two factors:

1) Some client operating systems make it impossible to override the DNS server settings obtained from DHCP. For instance, Android does not allow obtaining IP address assignments from DHCP while using static, user chosen DNS servers.

2) The gateway device that provides the DHCP server for the home is increasingly not a retail off-the-shelf appliance, but a device leased from the ISP that runs firmware approved by the ISP. Large ISPs in particular appear to be restricting or removing the ability for the subscriber to configure the DNS server settings in the gateway devices that they provide to customers. That is, this functionality *only* appears to be disappearing from ISP-branded hardware, not from retail gateway devices, which have retained this basic customization feature continuously since they first appeared on the market.

This presentation will mainly focus on the second factor, because even if all client devices supported manual DNS server configuration, it would be very tedious to statically configure every device in one's home, and tedium is a form of friction.

Are there workarounds for the technically savvy user? Yes, but there are tradeoffs in terms of extra expenses or technical quality. For instance, any technical workaround that requires purchasing additional hardware is a workaround that is unavailable to users who cannot afford that additional expense.

Why are some ISPs doing this? Are there arguably legitimate security reasons for this? Can/should these ISPs be convinced to stop in the interest of consumer choice?

Summary

Talk Duration

30 Minutes

Primary author(s) : EDMONDS, Robert (Fastly, Inc.)

Presenter(s) : EDMONDS, Robert (Fastly, Inc.)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **59** Type : **Standard Presentation**

Additional Truncated Response

Friday, 9 March 2018 14:00 (30)

Recently, L. Song, ⊠⊠⊠, of BII proposed the use of Additional Truncated Responses (draft-song-atr-large-resp-00) as way to improve resolution success rates for clients in the presence of large DNS responses.

We are using our large distributed measurement platform to evaluate the effect of the proposed behaviour by implementing a modified DNS server that implements ATR behaviour. This talk will present our findings.

Summary

Talk Duration

30 Minutes

Primary author(s) : Mr. HUSTON, Geoff (APNIC); Mr. SILVA DAMAS, Joao Luis (Bond Internet Systems)

Presenter(s) : Mr. HUSTON, Geoff (APNIC); Mr. SILVA DAMAS, Joao Luis (Bond Internet Systems)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **60** Type : **Standard Presentation**

KSK Sentinel

Friday, 9 March 2018 10:00 (30)

The KSK Roll is coming – but we still don't have good visibility into what the effects will be. KSK Sentinel (draft-ietf-dnsop-kskroll-sentinel) provides a way to measure what the **user** effect will be (and also allows mass measurement, using ads).

Summary

Talk Duration

30 Minutes

Primary author(s) : KUMARI, Warren (N/A)

Presenter(s) : KUMARI, Warren (N/A)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **61** Type : **Standard Presentation**

OARC Systems Update

Thursday, 8 March 2018 10:25 (20)

Summary

Talk Duration

15 Minutes

Primary author(s) : Mr. SOTOMAYOR, William (DNS-OARC)

Presenter(s) : Mr. MITCHELL, Keith (DNS-OARC)

Session Classification : OARC Business

Track Classification : OARC Business

Contribution ID : **62** Type : **not specified**

OARC Software Update

Thursday, 8 March 2018 10:05 (20)

Summary

Talk Duration

Presenter(s) : Mr. LUNDSTRÖM, Jerry (DNS-OARC)

Session Classification : OARC Business

Track Classification : OARC Business

Contribution ID : **63** Type : **not specified**

OARC Status Report

Thursday, 8 March 2018 09:35 (30)

Summary

Talk Duration

Presenter(s) : Mr. MITCHELL, Keith (DNS-OARC)

Session Classification : OARC Business

Track Classification : OARC Business

Contribution ID : **64** Type : **not specified**

OARC Governance Update/Discussion

Thursday, 8 March 2018 10:45 (15)

A chance to update OARC Members on governance developments with OARC Board representation since the AGM, and a chance for any feedback/discussion.

Summary

Talk Duration

Presenter(s) : WESSELS, Duane (Verisign); Mr. MITCHELL, Keith (DNS-OARC)

Session Classification : OARC Business

Track Classification : Members-Only

Contribution ID : **65** Type : **not specified**

Introduction from OARC Chairman

Thursday, 8 March 2018 09:30 (5)

Summary

Talk Duration

Presenter(s) : WESSELS, Duane (Verisign); Mr. MITCHELL, Keith (DNS-OARC)

Session Classification : OARC Business

Track Classification : OARC Business

Contribution ID : **66** Type : **not specified**

PGP Key Signing Session

Friday, 9 March 2018 13:15 (40)

Summary

Talk Duration

Presenter(s) : POUNSETT, Matthew

Contribution ID : **69** Type : **Standard Presentation**

Analyzing and Mitigating Privacy with the DNS Root Service

Friday, 9 March 2018 16:00 (30)

Processing of all DNS requests start at the root of the DNS tree and make use of either cached data from previous requests, or by traversing the DNS tree for the missing information. When \emph{QNAME minimization} is not in use, queries forwarded to the parental nodes in the DNS tree may leak private DNS query data. In this paper we examine 31 days during the month of January 2017 of queries sent from two recursive resolvers placed in two residential networks to the DNS root server operated by \isi's, analyzing the leaked QNAMEs for an impact on the network's privacy. We then compare a few DNS privacy preserving techniques against the privacy analysis against these networks.

Summary

Talk Duration

30 Minutes

Primary author(s) : HARDAKER, Wes (USC/ISI)

Presenter(s) : HARDAKER, Wes (USC/ISI)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **71** Type : **Lightning Talk**

gTLD Name Collisions 2012, next round

Friday, 9 March 2018 17:10 (10)

[See attached PPTX]

Summary

How to deal with name collisions caused by delegation of new TLDs is currently under discussion in the ICANN community. The talk summaries what happened in the 2012-round and what is being thought for future procedures.

Talk Duration

15 Minutes

Primary author(s) : Mr. KUHL, Rubens (NIC.br)

Presenter(s) : Mr. KUHL, Rubens (NIC.br)

Session Classification : Lightning Talks

Track Classification : Lightning Talks

Contribution ID : **72** Type : **Lightning Talk**

EDNS Compliance - Deprecating workarounds

Friday, 9 March 2018 17:20 (10)

The major open-source DNS server vendors has a plan to deprecate workarounds for broken EDNS implementations in servers. In this lightning talk, we (CZ.NIC, ISC, PowerDNS, NLnet Labs) will announce our plan to remove the workarounds from our DNS servers.

Summary

Talk Duration

Lightning

Primary author(s) : Mr. SURY, Ondrej (Internet Systems Consortium); Mr. VAN DIJK, Peter (PowerDNS); Mr. ŠPAČEK, Petr (CZ.NIC); DOLMANS, Ralph (NLnet Labs)

Presenter(s) : Mr. SURY, Ondrej (Internet Systems Consortium); Mr. ŠPAČEK, Petr (CZ.NIC); DOLMANS, Ralph (NLnet Labs)

Session Classification : Lightning Talks

Track Classification : Lightning Talks

Contribution ID : **73** Type : **Lightning Talk**

Processing DITL data quickly at USC/ISI

Friday, 9 March 2018 17:00 (10)

I can present either screenshots or a live demo about how we process B-Root specific DITL data with infrastructure at USC/ISI. I'll describe the format we store data in, how it enables us to rapidly perform analysis on it, and how we can do bulk processing of DNS requests after it has been converted to our textual saved format.

Summary

Talk Duration

Lightning

Primary author(s) : HARDAKER, Wes (USC/ISI)

Presenter(s) : HARDAKER, Wes (USC/ISI)

Session Classification : Lightning Talks

Track Classification : Lightning Talks

Contribution ID : **74** Type : **not specified**

Algorithm Rollover what is the safest approach

Friday, 9 March 2018 16:30 (10)

Summary

Talk Duration

Primary author(s) : Mr. DE CARVALHO NEVES, Frederico Augusto (Nic.br)

Presenter(s) : Mr. DE CARVALHO NEVES, Frederico Augusto (Nic.br)

Session Classification : Lightning Talks

Contribution ID : **75** Type : **Lightning Talk**

BIND testing with Meltdown Spectre

Friday, 9 March 2018 16:50 (10)

Results of testing different versions of BIND with Meltdown / Spectre patches installed on Redhat 7 servers.

Summary

Talk Duration

Lightning

Primary author(s) : DEVRIES, Peter (Quotient Inc)

Presenter(s) : DEVRIES, Peter (Quotient Inc)

Session Classification : Lightning Talks

Track Classification : Lightning Talks

Contribution ID : **76** Type : **not specified**

DITL Anonymization

Friday, 9 March 2018 16:40 (10)

Summary

Talk Duration

Primary author(s) : HOFFMAN, Paul (ICANN)

Presenter(s) : HOFFMAN, Paul (ICANN)

Session Classification : Lightning Talks

Track Classification : Lightning Talks