OARC 2015 Fall Workshop (Montreal)

Report of Contributions

Contribution ID : **0** Type : **not specified**

OARC President's Report

Saturday, 3 October 2015 10:05 (20)

Summary

Please also consider this submission for the NANOG65 DNS track

No

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

Session Classification : Members Session

Track Classification : OARC AGM

Contribution ID : **1** Type : **not specified**

OARC Systems Report

Saturday, 3 October 2015 11:15 (20)

Summary

Please also consider this submission for the NANOG65 DNS track

No

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

Session Classification : Members Session

Track Classification : Public Workshop

Contribution ID : **2** Type : **not specified**

DNS Software Test Centre Proposal

Saturday, 3 October 2015 11:55 (20)

Summary

Please also consider this submission for the NANOG65 DNS track

No

Primary author(s) : Mr. WEBER, Ralf (Nominum Inc)

Presenter(s) : Mr. WEBER, Ralf (Nominum Inc)

Session Classification : Members Session

Track Classification : Members-Only

Contribution ID : **4** Type : **not specified**

Impact of unknown EDNS options on the DNS

Saturday, 3 October 2015 16:30 (15)

The EDNS (Extension mechanisms for DNS) protocol allows us to add new features to DNS that were not envisioned when DNS was originally specified. DNSSEC, Client-subnet Identifier and DNS cookies are applications that use EDNS.

It appears from ISC's testing that a significant percentage of sites that support EDNS do not respond well to unknown EDNS options. The failure mode can be as severe as disabling EDNS (breaking DNSSEC). We are reluctant to encourage the use of new EDNS options until there is better tolerance for unknown EDNS options in the DNS. We would like to raise awareness of the issue, and find out what the community thinks we should do to address it.

This presentation will review the results of our testing and the current EDNS failure modes we see, and explain how to test your own site for compliance.

Summary

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : RISK, victoria (isc)

Presenter(s) : RISK, victoria (isc)

Session Classification : Public Workshop

Track Classification : Lightning Presentations

Contribution ID : **5** Type : **not specified**

Benchmarking of authoritative DNS servers and DNSSEC impact assessment

Saturday, 3 October 2015 16:00 (30)

CZ.NIC Labs created and continues to actively develop Knot DNS authoritative DNS server. The development team puts substantial effort into optimizing the server performance and searching for new optimization opportunities. So we created a DISTEL-based lab for benchmarking not only our server but for comparing many different authoritative DNS servers and versions.

The presentation shows our method for collecting data, explain statistics that we use for testing hypotheses about the server performance and presents results for Knot 2.0 and others with regard to mixed DNSSEC and non-DNSSEC traffic.

Summary

Development of Knot DNS authoritative server is supposed to be driven by benchmarks. Generally we want to test all changes that might affect performance and compare them to the previous versions. The basic question is whether we can see any statistically significant improvement, especially in case the changes in measurement results are small and unevenly distributed. To answer that question we use Hotelling's test and occasionally ANOVA and regression analysis to go a bit deeper and provide developers with information they are interested in. Another application of these methods is assessment of DNSSEC performance impact by comparing different DNSSEC algorithms on the same data sets and same servers.

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : Mr. HLAVACEK, Tomas (CZ.NIC, z.s.p.o.)

Presenter(s) : Mr. HLAVACEK, Tomas (CZ.NIC, z.s.p.o.)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **7** Type : **not specified**

A study of caching behavior with respect to root server TTLs

Sunday, 4 October 2015 14:00 (30)

The Root Server System Advisory Committee (RSSAC) within ICANN was recently tasked with considering the extent to which the current root zone TTLs are appropriate for today's Internet environment and the impacts of root TTL alterations on the wider DNS system. The historical DITL data from 2014 and 2015 was analyzed for trends in TTL adherence, answering some of the following questions: To what degree do root zone TTLs matter? Is there a difference in behavior for authoritative versus non-authoritative data? Do all TLDs exhibit similar inter-query time distributions? Do specific recursive implementations, ISPs, open resolvers, etc. diverge from general TTL adherence trends? How has inter-query time changed over the past two years? Would a change in root zone TTLs result in a change in traffic levels at root name servers?

Summary

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : WESSELS, Duane (Verisign); THOMAS, Matthew (Verisign)

Presenter(s) : WESSELS, Duane (Verisign); THOMAS, Matthew (Verisign)

Session Classification : Public Workshop: Root Data Analysis

Track Classification : Public Workshop

Contribution ID : **8** Type : **not specified**

Continuous Integration & Continuous Deployment - For the new nameserver infrastructures of DENIC eG

Sunday, 4 October 2015 09:30 (30)

This Abstract describes the concepts, the development and the functionalities of the DENIC DNS Countinuous Integration and Deployment Pipeline. Furthermore the advantages you could receive through this technics and automated testing. More informations will follow in the summary form.

Summary

Author: Christian Petrasch – DNS Operations / DENIC eG

Motivation:

The nameserver infrastructure of a TLD or TLD-like company is, concerning widespreading anycast infrastructures, a service where fast deployments of servers are a necessary feature. This applies to changes onto the server platform (software and OS and configurations) as well as reinstall or add new servers to the existing infrastructure. There are multiple reasons for this, like for example adding new customers,
mitigation of attacks, deploying new software or patches for every type of software (also the operating system) running on the platform.

Nevertheless, the stability and security of the servers shouldn't be influenced by increasing the deployment speed. Because of this, automated testing comes in the game. With an automated testing pipeline no new software or patch will be deployed in production without a successful integration test. This guarantees that
the stability of the running production will be provided after the deployment of the new feature.

The orchestration is the next step in this structure. Orchestration melts fast deployment, testing and controlling the service together. The orchestration is a toolstack for controlling the dependencies between new deployments and a maybe unavoidable outtake of a server during reinstallation and the absolutely uninterrupted service for the customer. This is the reason why an orchestration is the mandatory connection between controlling the routing equipment and the server equipment.

The presentation should give a deep overview about the development and the techniques DENIC is using for building and testing the DNS infrastructure and the big wins you can get because of Continuous Delivery in a DNS platform. This includes a comparison between the old infrastructure and the new one, operations, problem handling (debugging, attack mitigation), usability and orchestration.

Furthermore it should show the advantages of CI/CD and automated testing at nameserver structures for implementing new updates, software or any changes which can be done on servers.

DNS Pipeline Infrastructure

Structure of the presentation:

- Overview Old infrastructure
- used techniques
- Problems with old techniques
- Short overview about what is Continuous Integration & Continuous Deployment

- How to reach the goal – Development Steps to CI/CD
- Developing concepts (what is a good way and what not ?)
  - which type of configuration management
  - which type of authoritative datasource for the CMDB
- Developing the necessary processes
  - testing processes/pipeline
  - deployment strategy pipeline
- Decision of tools (Virtualization, Automation, GUI, Testing, Routing Control)
- Steps
  - Authoritative datasource
  - Build necessary toolstack
  - Building a staging pipeline
  - Orchestration
  - Steering the routing (Bonus)
- Benefits after reaching the goal
  - Deployment speed (Live DEMO)
  - Rolling out new servers
  - Reliability, (Consistency, better ways to update)
  - Automated testing
  - Mitigation of attacks
  - Implementation of monitoring

Estimated duration: approximately 30 min

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : Mr. PETRASCH, Christian (DENIC eG)

Presenter(s) : Mr. PETRASCH, Christian (DENIC eG)

Session Classification : Public Workshop: Operations Track

Track Classification : Public Workshop

Contribution ID : **10** Type : **not specified**

Thirteen Years of ”Old J Root”

Sunday, 4 October 2015 15:00 (30)

Thirteen years ago Verisign renumbered j.root-servers.net so that it could be anycasted. Since that time, we have been continuing to answer queries sent to the old IP address. We have also been collecting some data on queries to old J-root.

In this presentation we will explore such questions as: what do we know about the clients of old J-root? Do they overlap with clients of the real J-root? Are there noticeable differences in traffic characteristics (e.g., EDNS, DNSSEC, query types) between the two? Does old J-root traffic fluctuate in the same way as real traffic? When real J-root gets attacked, does old J-root also get attacked? If so, can this be used to identify attacks coming through recursive name servers?

Summary

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : WESSELS, Duane (Verisign)

Co-author(s) : Mr. CASTONGUAY, Jason (Verisign); Mr. BARBER, Piet (Verisign)

Presenter(s) : WESSELS, Duane (Verisign)

Session Classification : Public Workshop: Root Data Analysis

Track Classification : Public Workshop

Contribution ID : **11** Type : **not specified**

Happy DNS Eyeballs?

Sunday, 4 October 2015 17:15 (30)

Much work has been undertaken in the browser world to produce the so-called "Happy Eyeballs" outcome. This is an outcome where the client will detect if the service is a dual stack service and if so then use a connection process that slightly biases the client in favour of using IPv6 as the transport for the DNS. What evidence is there for a similar mode of behaviour of DNS resolvers? This presentation will report on a large scale measurement experiment that was intended to expose the protocol behaviour of resolvers and determine whether they have any protocol selection bias.

Summary

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : Mr. HUSTON, Geoff (APNIC)

Co-author(s) : MICHAELSON, George (APNIC P/L)

Presenter(s) : Mr. HUSTON, Geoff (APNIC)

Session Classification : Public Workshop: Resolvers Track

Track Classification : Public Workshop

Contribution ID : **13** Type : **not specified**

OpenDNS; Managing DDoS Attacks

Sunday, 4 October 2015 09:00 (30)

Open resolvers will always be a target for abuse either as an attack amplification point or as a mask of the attack source. This presentation discusses the measures that OpenDNS has put in place to ensure that their open resolvers contribute towards reducing or blocking DDoS attacks. It goes on to discuss future plans to identify limit or block DoS sources.

Summary

Indico rendering error

```
Could not include image: [404] Error fetching image
```

This talk discusses OpenDNS' four main DDoS preventative measures:

- Rate Limiting
- The Droplist
- The Freezelist
- Authoritative RTT Handling

It also explores some possible future technologies:

- TClist
- Freezelist Thawing
- NXDOMAIN per-level counting
- Whitelist Labels

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : Mr. SOMERS, Brian (OpenDNS, FreeBSD)

Presenter(s) : Mr. SOMERS, Brian (OpenDNS, FreeBSD)

Session Classification : Public Workshop: Operations Track

Track Classification : Public Workshop

Contribution ID : **14** Type : **not specified**

Publishing zone scan data using an open data portal

Sunday, 4 October 2015 12:15 (15)

NZRS has been running zone scans on a monthly basis over the active .nz domain names for the past two years. We are using dnscheck with custom changes to collect DNS health information, as well as IPv6 deployment metrics. The data is of no use if it can't be made readily available to interested parties.

To sort this out, NZRS will start using an open data portal provided by Socrata to allow open access to the zone scan and other datasets about the registry operations. The data portal will allow to download the data, but also to explore it in a visual way.

This presentation will cover
- Methodology and infrastructure to run the zone scan
- Overview of the data collected
- Introduction to the open data portal (possibly a demo)
- Some interesting findings.

Time required: Ideally 30 minutes, but can be adjusted to be a lightning talk.

Summary

Please also consider this submission for the NANOG65 DNS track

No

Primary author(s) : Mr. CASTRO, Sebastian (NZRS)

Presenter(s) : Mr. CASTRO, Sebastian (NZRS)

Session Classification : Public Workshop: Data Analysis

Track Classification : Public Workshop

Contribution ID : **15** Type : **not specified**

Analyzing the distribution of DNS clients to recursive name servers across the Internet

Sunday, 4 October 2015 16:00 (30)

As a byproduct of our web-based Real User Monitoring (RUM) agent, Dyn obtains the IP addresses of individual hosts running web browsers all over the world as well as the IP addresses of the recursive servers queried by those hosts. We have analyzed a rich data set of over 110 million client IP address-to-recursive IP address mappings to research an area of DNS that we believe has not been sufficiently studied. For example, what is the distribution of the number of clients per recursive server? Where are clients located relative to the recursive servers they use, both from a geographic as well as network topological perspective? What query patterns do individual clients follow if they use multiple recursive servers? We report on these and other interesting findings.

Summary

Please also consider this submission for the NANOG65 DNS track

No

Primary author(s) : LARSON, Matt (Dyn, Inc.)

Presenter(s) : LARSON, Matt (Dyn, Inc.)

Session Classification : Public Workshop: Resolvers Track

Track Classification : Public Workshop

Contribution ID : **16** Type : **not specified**

Neutering ANY queries: how we did it

Sunday, 4 October 2015 10:00 (30)

DNS ANY queries are a source of controversy and strong feelings. In practice ANY queries are used for debug purposes, but are frequently abused in amplification attacks, as they give the best amplification factor. In some non-traditional DNS authoritative servers the cost of answering ANY queries can be high due to multiple DB lookups and dynamic records.
Once in a while someone thinks that using ANY query is a good way to reliably get all RRsets in one query, frequently without understanding the semantics or implications.
We have explored a number of alternatives to answer ANY queries without breaking any deployed systems, and at the same time discourage the use of ANY query.
In this talk we will cover the alternatives and present our solution to a cacheable, non-breaking "negative" answer to ANY queries.

Summary

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : Mr. GUÐMUNDSSON, Ólafur (CloudFlare Inc.)

Co-author(s) : Mr. VALSORDA, Filippo (CloudFlare); Mr. MAJKOWSKI, Marek (CloudFlare)

Presenter(s) : Mr. GUÐMUNDSSON, Ólafur (CloudFlare Inc.)

Session Classification : Public Workshop: Operations Track

Track Classification : Public Workshop

Contribution ID : **17** Type : **not specified**

Internet Performance Impacts of Canadian Content Hosting

Sunday, 4 October 2015 11:00 (30)

In addition to driving ccTLD growth, a strong national content hosting industry improves consumer experience by making content faster, cheaper, and more accessible. Ironically, the existence of a large content-hosting industry next door has tended to artificially reduce the percentage of both Canadian and Mexican content that is domestically hosted.

In this talk, we'll examine the most popular domains in Canada (using Alexa ranking data), as well as the much broader spectrum of domains hosted in .CA. We'll utilize BGP routing and IP geolocation to assess the geographies and key providers that support Canada's current state of domestic versus international content placement. Beyond this basic census, we'll also examine some of the potential performance impacts of nonlocal hosting on content consumers in both Eastern and Western Canada.

Summary

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : COWIE, Jim (Dyn, Inc.); LARSON, Matt (Dyn, Inc.)

Presenter(s) : COWIE, Jim (Dyn, Inc.); LARSON, Matt (Dyn, Inc.)

Session Classification : Public Workshop: Data Analysis

Track Classification : Public Workshop

Contribution ID : **18** Type : **not specified**

DNS big data analytics

Sunday, 4 October 2015 11:30 (30)

Introduction

As the operator of the .nl ccTLD, SIDN is very interested in keeping the .nl zone as safe as possible.
Analyzing the query data can help to detect cybercrime activity in the .nl zone which we can than try to cleanup.
Traditional DNS query data analysis done by storing data as PCAP's and analyzing them with tools such a tshark and wireshark is often a slow and painful process.
When storing DNS query data as PCAP files makes you will quickly run into performance and scalability problems.
Most tooling used to analyze PCAP's is single threaded and has limited or no sql compatibility.
What is required is a system which can cope with large volumes of PCAP data and still offer good query performance.

That's why SIDN developed a DNS big data platform called ENTRADA, this platform is built on top of the Hadoop stack using open source technology.
DNS query data from our authoritative name servers is stored on this platform and can be analyzed using multiple interfaces and languages.
The system supports SQL, which means that anyone with SQL knowledge can quickly start analyzing the query data.
Currently the database contains over 64 billion DNS queries, each day some 130 million new queries are added and this number will grow as we hook up more name servers.

In this presentation I will be talking about system design, use cases and our experiences.

Platform design

The platform at SIDN is used by the R&D team and is quite small (5 nodes)
The costs of setting up such a cluster are very modest, the main components are as expected hardware and people.
The hardware does not have to be enterprise grade and much of the required knowledge is available for free online.

Adding more storage and compute capacity is as simple as adding more disk drives or servers.
The cluster storage capacity at the moment is about 100 billion DNS queries and this data can be queried very efficiently. Depending on the type
of query and number of data partitions that have to be scanned, most queries will return a result within seconds.

Privacy

Privacy is an important aspect when collecting DNS data, because the query data might reveal personal information about the users who are sending DNS queries.
The IP address of a client can in some cases be used to identify and track users (for a home user operating a private resolver, or for small shared resolvers)
We designed a novel privacy framework (1) because it introduces privacy management to the use of DNS data
and (2) because, to that end, it integrates legal, organizational and technical aspects of privacy management.

This is described in our paper: https://www.sidnlabs.nl/uploads/tx_sidnpublications/SIDN_Labs_Privacyraamwerk_Positio

Workflow

The time it takes from a query being received on the name server until it is available in the database for analysis is just a couple of minutes.
The steps involved are:

- get pcap data every x minutes from NS
- PCAP conversion
- enrichment of data
- storage
- query!

Storage

There are a lot of different storage technologies, we chose to use the Parquet format to encode the data and Hadoop HDFS as a distributed storage layer.
This part explains why Parquet is such a good fit for storing DNS data.

- Why we chose Parquet
- Size difference (pcap vs parquet, total database size)
- How do you convert pcap data to parquet (write parquet with Avro schema (KiteSDK))
- Parquet format can be read by Impala but also by Spark, this makes it very flexible.

Query engines and interfaces

The data stored in the system can be access through multiple query engines and interfaces.
The support workloads from a simple sql query to advanced graph and machine learning jobs.

Impala/Impyla (SQL engine)
Spark (SQL/Graph/Machine learning engine)
Hue (SQL web interface)
Jupyter (python notebook)

use cases

Focused on increasing the security and stability of .nl

- DNS security App (visualize traffic patterns for phishing domain names)
- Botnet detector (detect botnet infections and report these to abuse information exchange (https://www.abuseinformationexchange.nl/english))
- Real-time Phishing domain name detection
- Statistics dashboard (stats.sidnlabs.nl)
- Scientific research (collaboration with Dutch Universities)
- Ad-hoc operational analysis (quick analysis of current issues in the DNS)

experiences

Our experiences in working with this data:
So much work to be done when this data is available, we hired an additional Data scientist.

Future work:
- Combine passive data from .nl authoritative name server with active scans of the complete .nl zone and ISP data.
- Adding more name servers and resolvers.
- Open data interface

Summary

1. We believe that our choice of technology combined with our privacy framework is quite novel.
2. Our setup proves that a big data platform can start small with limited costs and still be powerful.
3. We provide the rational behind our architectural decisions with regards to tools, workflow and data formats for storage.
4. We provide example use cases of what is possible when this data is available for analysis.

Summary

As the operator of the .nl ccTLD, SIDN is very interested in keeping the .nl zone as safe as possible. For this goal a DNS big data platform (ENTRADA) has been developed by SIDN.
This talk will provide an overview of the ENTRADA DNS big data platform design, technology and use cases.

This talk with provide insight into the following aspects of DNS big data

1. Big data does not have to mean big costs. It's possible to build a DNS big data cluster with low budget and make it grow. 2 What does such a platform look like? What are the architectural decisions? (tools, workflow, data formats for storage) How do you address privacy concerns?
2. What are the use cases for DNS big data? We have developed some applications and are doing research together with universities.

The focus will be on:
a) Major system components
b) The technology used to build the platform.
c) the possible use cases for DNS data analytics such as botnet detection and anti phishing.

keywords: DNS, Hadoop, Apache Parquet, Cloudera Impala/Impyla, Jupyter, Apache Spark.

Please also consider this submission for the NANOG65 DNS track

No

Primary author(s) : Mr. WULLINK, M (SIDN)

Presenter(s) : Mr. WULLINK, M (SIDN)

Session Classification : Public Workshop: Data Analysis

Track Classification : Public Workshop

Contribution ID : **19** Type : **not specified**

Real World Impacts of EDNS Client Subnet

Sunday, 4 October 2015 16:30 (30)

Client Subnet offers the ability to offer better geolocation of end-users via DNS responses. This talk will concentrate on what happens when Client Subnet is enabled on a public resolver. It will look at upstream traffic patterns, cache performance, and other factors that come into play with Client Subnet. At the end of this talk, DNS providers should have a better idea of how Client Subnet will impact their performance & network.

Summary

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : HARTVIGSEN, Brian (OpenDNS)

Presenter(s) : HARTVIGSEN, Brian (OpenDNS)

Session Classification : Public Workshop: Resolvers Track

Track Classification : Public Workshop

Contribution ID : **20** Type : **not specified**

Cluster the long tailed domains base on passiveDNS.cn

Sunday, 4 October 2015 12:00 (15)

Cluster the DNS domains is a basic but very important work in analyzing the dizzy businesses of the Internet. Only based on the accurate clustered domain result, we can discern and analyze all kinds of DNS data. Now, most of the works focus on the domain structure and hoping finding the relationships among kinds of domains. Recently, based on the largest public passiveDNS database in China, we are exploring some new but beneficial ways on cluster the long tailed domains(based on some filter rules). Except the domain structure, we add two dimensions: client and server data. Introduce the real data of up-down stream is a big extension, of course it's more accurate. From the test result, the two dimensions is helpful in clustering the domains and finding the both benign and malicous domain communities.

Summary

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : Mr. ZHANG, zaifeng (QIHOO 360)

Co-author(s) : Mr. ZHAO, zhihui (QIHOO 360)

Presenter(s) : Mr. ZHANG, zaifeng (QIHOO 360)

Session Classification : Public Workshop: Data Analysis

Track Classification : Lightning Presentations

Contribution ID : **22** Type : **not specified**

F-root Anycast Research using RIPE Atlas

Sunday, 4 October 2015 14:30 (30)

ISC has been using data routinely collected by every RIPE Atlas node to research the effectiveness of F-root's current transit and peering arrangements.

The presentation will show how visualisation of this data can identify issues that should be resolved, along with "before and after" pictures showing the effect of changes that we already made to our routing configuration based on this analysis.

Summary

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : Mr. BELLIS, Ray (Internet Systems Consortium, Inc.)

Presenter(s) : Mr. BELLIS, Ray (Internet Systems Consortium, Inc.)

Session Classification : Public Workshop: Root Data Analysis

Track Classification : Public Workshop

Contribution ID : **23** Type : **not specified**

Next Steps in DANE Adoption

Saturday, 3 October 2015 15:00 (30)

This talk will discuss upcoming and future steps envisioned to increase the adoption of DNSSEC and DANE (DNS-based Authentication of Named Entities) by Internet applications. It will start by providing an overview of the current state of adoption of DANE, and then discuss challenges faced by some application communities and some applications for which DANE doesn't yet provide a solution. Among the topics covered will be a proposed new TLS extension to allow servers to deliver a DANE record and the associated DNSSEC chain to clients, a mechanism to allow the use of TLSA records for client authentication, and others.

Summary

This talk will discuss upcoming and future steps envisioned to increase the adoption of DNSSEC and DANE (DNS-based Authentication of Named Entities) by Internet applications. It will start by providing an overview of the current state of adoption of DANE, and then discuss challenges faced by some application communities and some applications for which DANE doesn't yet provide a solution. Among the topics covered will be a proposed new TLS extension to allow servers to deliver a DANE record and the associated DNSSEC chain to clients, a mechanism to allow the use of TLSA records for client authentication, and others.

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : Mr. HUQUE, Shumon (Verisign Labs)

Presenter(s) : Mr. HUQUE, Shumon (Verisign Labs)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **24** Type : **not specified**

An Overview of DNS Privacy Mechanisms

Saturday, 3 October 2015 14:00 (30)

As part of the IETF's efforts to secure all protocols against pervasive surveillance, several privacy enhancements to the DNS are actively being developed with prototype implementations of such enhancements also emerging. This talk will provide a technical description of these mechanisms as well as deployment challenges and related considerations. Topics to be covered will include query-name minimization, DNS over TLS/DTLS and other encryption proposals, payload padding, etc. We will also cover related efforts to minimize leakage of DNS names in other protocols.

Summary

As part of the IETF's efforts to secure all protocols against pervasive surveillance, several privacy enhancements to the DNS are actively being developed with prototype implementations of such enhancements also emerging. This talk will provide a technical description of these mechanisms as well as deployment challenges and related considerations. Topics to be covered will include query-name minimization, DNS over TLS/DTLS and other encryption proposals, payload padding, etc. We will also cover related efforts to minimize leakage of DNS names in other protocols.

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : MANKIN, Allison (Verisign Labs)

Co-author(s) : Mr. HUQUE, Shumon (Verisign Labs)

Presenter(s) : MANKIN, Allison (Verisign Labs); Mr. HUQUE, Shumon (Verisign Labs)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **25** Type : **not specified**

Using TLS for DNS privacy in practice

Saturday, 3 October 2015 14:30 (30)

This talk will focus on the existing prototype implementations of DNS-over-TLS and dive into some of the finer points of using TLS in
practice. This will include authentication issues, performance considerations, TCP connection management, DoS mitigation and a demonstration. It will also discuss the current best practices for using TLS in applications and the upcoming developments in TLS 1.3.

Summary

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : DICKINSON, Sara (Sinodun IT)

Presenter(s) : DICKINSON, Sara (Sinodun IT)

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **26** Type : **not specified**

dnstap-whoami: one-legged exfiltration of resolver queries

Sunday, 4 October 2015 17:00 (15)

A few existing "whoami" or "dnsecho" authoritative DNS services allow for limited extraction of information about the resolver to the original client that would normally be hidden. For example, querying an anycasted resolver like 8.8.8.8 with the command "dig @8.8.8.8 whoami.akamai.net" will return an address record revealing a unicast initiator address used by the anycast service. This is "one-legged", because the original client only has visibility into the stub/recursive "leg" of the DNS interaction. The DNS-OARC porttest tool is another example of a "one-legged" service.

Similarly, many DNS research projects use special purpose zones with instrumented nameservers which capture incoming query packets for analysis. For example, scans for open recursive DNS servers typically control both the stub/recursive "leg" and the recursive/authoritative "leg" and are thus "two-legged". This requires a more heavyweight investment but results in a richer set of data.

This talk will demonstrate an enhanced "whoami" authoritative DNS server that can exfiltrate more detailed information about the recursive/authoritative interaction to the original client, including the complete resolver query packet sent to the authoritative server, using the dnstap format to compactly tunnel structured information which can be decoded by the original client.

Summary

Please also consider this submission for the NANOG65 DNS track

Primary author(s) : EDMONDS, Robert (Farsight Security, Inc.)

Presenter(s) : EDMONDS, Robert (Farsight Security, Inc.)

Session Classification : Public Workshop: Resolvers Track

Track Classification : Lightning Presentations

Contribution ID : **27** Type : **not specified**

Benchmarking and profiling DNS systems with modern Linux tools

Saturday, 3 October 2015 16:45 (15)

This talk will outline the use of tools from the netsniff-ng toolkit and the Linux kernel along with a home-grown benchmark harness to characterize UDP DNS performance. These tools operate very differently from "traditional" utilities like dnsperf/resperf and produce very different results, possibly contradicting conventional wisdom that UDP on Linux is slow.

Summary

Please also consider this submission for the NANOG65 DNS track

No

Primary author(s) : EDMONDS, Robert (Farsight Security, Inc.)

Presenter(s) : EDMONDS, Robert (Farsight Security, Inc.)

Session Classification : Public Workshop

Track Classification : Lightning Presentations

Contribution ID : **28** Type : **not specified**

Impact of DNS over TCP - a resolver point of view

Saturday, 3 October 2015 17:00 (30)

Using traffic captured at two different ISP's recursive resolvers we analise the potential impact on the servers of long lived TCP sessions, investigating the effect of timeout settings, the total number of simultaneous connections that would be kept open and the potential benefits of connection reuse as proposed in the current version of draft-ietf-dnsop-5966bis, with the intent of offering simulated operational advice, based on observerd traffic.

The study looks at the impact on the recursive server as it queries authoritative servers as well as while it talks to stubs, two very different aspects of the life of a recursive server.

Summary

Please also consider this submission for the NANOG65 DNS track

Yes

Primary author(s) : Mr. SILVA DAMAS, Joao Luis (Bond Internet Systems); Dr. CORTÉS SACK, Mónica (DIT-UPM/Bond Internet Systems)

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

Session Classification : Public Workshop

Track Classification : Public Workshop

Contribution ID : **29** Type : **not specified**

Idea: DNS over QUIC / zone transfer over QUIC or TLS/TCP

Saturday, 3 October 2015 17:30 (15)

The presentation discusses just an idea
about DNS over QUIC and zone transfer
over QUIC or TLS/TCP.
The third transport of DNS may be QUIC.
Both DNS and QUIC use UDP and port 53/UDP may be possible to share.
(If possible, implementation status will be reported, but it seems hard.)
And zone transfers may be performed over QUIC or TLS/TCP transport with
server certificate authentication.

Summary

Please also consider this submission for the NANOG65 DNS track

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 : Lightning Presentations

Contribution ID : **34** Type : **not specified**

OARC Treasurer's Report

Saturday, 3 October 2015 10:25 (10)

Summary

Please also consider this submission for the NANOG65 DNS track

No

Presenter(s) : WESSELS, Duane (Verisign)

Session Classification : Members Session

Track Classification : OARC AGM

Contribution ID : **35** Type : **not specified**

OARC Projects Update

Saturday, 3 October 2015 11:35 (20)

Summary

Please also consider this submission for the NANOG65 DNS track

No

Presenter(s) : Ms. KHEMLANI, Dalini (DNS-OARC)

Session Classification : Members Session

Track Classification : Public Workshop

Contribution ID : **36** Type : **not specified**

OARC Board Elections

Saturday, 3 October 2015 10:35 (10)

Summary

Please also consider this submission for the NANOG65 DNS track

No

Session Classification : Members Session

Track Classification : OARC AGM

Contribution ID : **37** Type : **not specified**

Introduction from OARC Chairman

Saturday, 3 October 2015 10:00 (5)

Summary

Please also consider this submission for the NANOG65 DNS track

No

Presenter(s) : Mr. FILIP, Ondrej (CZ.NIC)

Session Classification : Members Session

Track Classification : OARC AGM

Contribution ID : **38** Type : **not specified**

PGP Signing Session

Sunday, 4 October 2015 13:00 (60)

Summary

Please also consider this submission for the NANOG65 DNS track

Presenter(s) : Mr. VERGARA ERECHE, Mauricio (ICANN)

Contribution ID : **39** Type : **not specified**

Impact of unknown EDNS options on the DNS

Monday, 5 October 2015 16:25 (10)

Summary

Please also consider this submission for the NANOG65 DNS track

Presenter(s) : WINSTEAD, Eddy; RISK, victoria (isc)

Session Classification : NANOG65 DNS Track

Track Classification : NANOG65 DNS Track

Contribution ID : **40** Type : **not specified**

Benchmarking of authoritative DNS servers and DNSSEC impact assessment

Monday, 5 October 2015 16:35 (30)

Summary

Please also consider this submission for the NANOG65 DNS track

Presenter(s) : Mr. HLAVACEK, Tomas (CZ.NIC, z.s.p.o.)

Session Classification : NANOG65 DNS Track

Track Classification : NANOG65 DNS Track

Contribution ID : **41** Type : **not specified**

Managing DDoS Attacks

Monday, 5 October 2015 15:25 (30)

Summary

Please also consider this submission for the NANOG65 DNS track

Presenter(s) : Mr. SOMERS, Brian (OpenDNS, FreeBSD)

Session Classification : NANOG65 DNS Track

Track Classification : NANOG65 DNS Track

Contribution ID : **42** Type : **not specified**

F-root Anycast Research using RIPE Atlas

Monday, 5 October 2015 15:55 (30)

Summary

Please also consider this submission for the NANOG65 DNS track

Presenter(s) : Mr. BELLIS, Ray (Internet Systems Consortium, Inc.)

Session Classification : NANOG65 DNS Track

Track Classification : NANOG65 DNS Track

Contribution ID : **43** Type : **not specified**

DNS-OARC Overview

Monday, 5 October 2015 15:15 (10)

Summary

Please also consider this submission for the NANOG65 DNS track

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

Session Classification : NANOG65 DNS Track

Track Classification : NANOG65 DNS Track

Contribution ID : **44** Type : **not specified**

Root KSK Rollover

Monday, 5 October 2015 17:05 (10)

Summary

Please also consider this submission for the NANOG65 DNS track

Presenter(s) : AKPLOGAN, Adiel (ICANN)

Session Classification : NANOG65 DNS Track

Track Classification : NANOG65 DNS Track