OARC 2014 Spring Workshop (Warsaw)

DNS-OARC

Domain Name System Operations Analysis and Research Center

Report of Contributions

https://indico.dns-oarc.net/e/workshop-2014-05

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

IETF work on DNS privacy

Saturday, 10 May 2014 14:30 (30)

At the IETF 88 meeting in Vancouver, the first one which took into account the Snowden revelations, there was a lot of enthusiasm on action to improve the privacy on the Internet http://www.ietf.org/blog/2013/11/streng
the-internet/. This was summarized in a press release http://www.ietf.org/media/2013-11-07-internet-privacy-and-security.html claiming that "all of the working groups that considered the topic have started planning to address the threat using IETF tools that can mitigate aspects of the problem". Now, what is actually done in the DNS field?

Summary

There are several actions already under way:

- work on a future RFC on "DNS privacy problem statement". Two Internet-Drafts cover this. The work started in the perpass working group and will soon formally move to dnsop.
- work on "QNAME minimization", a simple and deployable technique to minimize the amount of data sent to authoritative name servers.
- work on encryption of DNS traffic. There are existing solutions (IPsec, DNScrypt) and possible new techniques (an Internet-Draft suggests a new method).

This talk will describe these actions, their current state after IETF 89 in London and the discussions they trigger.

Primary author(s) : Mr. BORTZMEYER, Stéphane (AFNIC)

Presenter(s) : Mr. BORTZMEYER, Stéphane (AFNIC)

Track Classification : Public Workshop

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

No Help Desk for Light Switches

Sunday, 11 May 2014 10:30 (30)

Increasing numbers of Internet-connected fridges and grandparents, together with cloud-based service delivery hysteria, are pushing availability requirements for web-accessible services through the roof. Subscribers are less interested in the reasons for failure, and are largely disinclined to try and call anybody for help (who would they call?) Service unavailability leads to lost subscribers, lost momentum and fear of lost investment and business failure. Being up is important.

Small, upstart web properties have options for outsourcing pieces of their infrastructure and operations to get a leg up on network and platform availability. With escalating availability requirements and a desire to be able to serve hot markets opportunistically, we consider how deep we can dig this particular rabbit-hole. We describe some of our thinking about how to scale our current service delivery platform from 20 sites globally to something much, much, much (much) bigger. We consider logistics, security, provisioning, manageability, monitoring and measurement, and begin to paint a picture of DNS service at a scale not previously seen on the Internet.

Summary

Primary author(s) : Mr. ABLEY, Joseph (Dyn, Inc.)

Presenter(s) : Mr. ABLEY, Joseph (Dyn, Inc.)

Contribution ID : **3** Type : **not specified**

A survey of Namecoin the peer to peer based DNS system

Sunday, 11 May 2014 14:10 (30)

Along side all of the new TLDs which have come into being, there is a dark horse: .bit. .bit isn't one of ICANN's most recently blessed TLDs, such as .guru, .democrat or .sexy, it is the top level domain which is served by the Namecoin infrastructure. The Namecoin platform seeks to provide an alternative (read as non-ICANN regulated TLD) decentralized domain name system built on a modified version of the Bitcoin software. The research summarized in this presentation is a review of the concepts behind Namecoin, its implementation, and the use of the Namecoin platform.
The first area of research was recreating the .bit zone and analyzing its contents. Specifically, it was focused on identifying the number of IP addresses associated with .bit domains, the autonomous systems which these IPs represent, and the type of records which appear most prominently. The next step was to acquire Namecoin and use it to create a domain with a collection of resource records. The goal of this exercise was to determine the barriers to entry and explore the user experience of the platform. The results were translated to a consideration for the usability of the .bit name space and its requirement for a copy of the Namecoin blockchain, a recursive DNS server with the .bit zone, or a web browser plugin. Namecoin has hurdles to overcome in the form of accessibility, integration, and hardening of the source. Its future is a race between innovation which can be seen in the development of DNSChain and its role as the new malware safe haven.

Summary

Until recent turbulence in the market, it was looking like it might be the year of the Bitcoin. Overstock.com started accepting it as legal tender and Google began to explore this option as well. In the shadow of Bitcoin rose a collection of alternative crypto currencies: Litecoin, Dogecoin, CoinyeWest, Hobonickles … the list goes on. Then, using the same ideas and some of the same codebase, came Namecoin, which acts as a decentralized domain name system for the global top level domain ".bit". Similar to ".gnu", ".zkey", ".onion", ".exit", and ".i2p" the use of the namespace isn't ICANN approved and is mentioned in the IETF Special-Use Domain Names of Peer-to-Peer Systems memo.

This presentation will start with an overview of Namecoin followed by a partial analysis of the contents of the .bit zone.
- The mechanics of Namecoin domain creation, record addition, and credit generation.
- Details of the block chain and how it provides a full transactional history of the namespace
- Zone contents analysis - Due to the decentralized nature of namecoin, generating a .bit zone file is a data extraction and transformation exercise.
- What types of records are most popular in this namespace? A vs. AAAA
- What is the distribution of IPs? Are there signs of domain squatting?
- What countries have the highest adoption rate?
- Implementation - How has the namespace has been abused? (the Necurs root kit) and How is it being used to fuel innovation? (DNSChain)

Primary author(s) : Mr. BAKER, Christopher (Dyn)

Presenter(s) : Mr. BAKER, Christopher (Dyn)

Track Classification : Public Workshop

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

getdns-api implementation

Sunday, 11 May 2014 09:00 (30)

Verisign and NLnet Labs have recently announced the first beta release (0.1.0) of an open source implementation of the getdns API specification. The project's home page is at http://getdnsapi.net.

getdns is a modern asynchronous DNS API. It implements DNS entry points from a design developed and vetted by application developers, in the specification at http://www.vpnc.org/getdns-api/ edited by Paul Hoffman.

With the implementation of this API, we intend to offer application developers a modernized and flexible way to access DNS security (DNSSEC) and other powerful new DNS features; a particular hope is to inspire application developers towards innovative security solutions in their applications.

In this presentation I will give an application developers view of DNSSEC and describe the independently written getDNS API specification. I will showcase the open source implementation of the specification built by our team of developers from NLNet Labs and Verisign.

The presentation will cover
how to perform resolution in all the different forms the different ways to perform DNSSEC and
the different levels of security assurances applications can get
the asynchronous support and how our implementation can integrate in
the application developers event base of choice the extensibility of the library
the limits of our current implementation and the roadmap for near-future development

Summary

Primary author(s) : Mr. TOOROP, Willem (NLnet Labs)

Presenter(s) : Mr. TOOROP, Willem (NLnet Labs)

Track Classification : Public Workshop

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

T-DNS: Connection-Oriented DNS to Improve Privacy and Security

Saturday, 10 May 2014 15:00 (30)

This talk will discuss *connection-oriented DNS* to improve DNS security and privacy. DNS is the canonical example of a connectionless, single packet, request/response protocol, with UDP as its dominant transport. Yet DNS today is challenged by eavesdropping that compromises privacy, source-address spoofing that results in denial-of-service (DoS) attacks on the server and third parties, injection attacks that exploit fragmentation, and size limitations that constrain policy and operational choices. We propose *t-DNS* to address these problems: it uses TCP to smoothly support large payloads and mitigate spoofing and amplification for DoS. T-DNS uses transport-layer security (TLS) to provide privacy from users to their DNS resolvers and optionally to authoritative servers.

Summary

Traditional wisdom is that connection setup will balloon latency for clients and overwhelm servers. We provide data to show that these assumptions are overblown—our model of end-to-end latency shows *TLS to the recursive resolver is only about 21% slower*, with UDP to the authoritative server. End-to-end latency is 90% slower with TLS to recursive and TCP to authoritative. Experiments behind these models show that after connection establishment, TCP and TLS latency is equivalent to UDP. Using diverse trace data we show that frequent connection reuse is possible (60−95% for stub and recursive resolvers, although half that for authoritative servers). With conservative timeouts (20 s at authoritative servers and 60 s elsewhere) we show that *server memory requirements match current hardware*: a large recursive resolver may have 25k active connections consuming about 9 GB of RAM. These results depend on specific design and implementation decisions—query pipelining, out-of-order responses, TLS connection resumption, and plausible timeouts.

We hope to solicit feedback from the OARC community about this work to understand design and operational concerns if T-DNS deployment was widespread. The work in the talk is Liang Zhu, Zi Hu, and John Heidemann (all of USC/ISI), Duane Wessels and Allison Mankin (both of Verisign), and Nikita Somaiya (USC/ISI).

A technical report describing the work is at
http://www.isi.edu/~johnh/PAPERS/Zhu14a.pdf

and the protocol changes are described as
http://datatracker.ietf.org/doc/draft-hzhwm-start-tls-for-dns/

Primary author(s) : Mr. HEIDEMANN, John (USC/Information Sciences Institute)

Presenter(s) : Mr. HEIDEMANN, John (USC/Information Sciences Institute)

Track Classification : Public Workshop

Contribution ID : **12** Type : **not specified**

Detecting and Clustering Botnet Domains Using DNS Traffic

Sunday, 11 May 2014 14:40 (30)

In this paper we focus on detecting and clustering distinct groupings of domain names that are queried by numerous sets of infected machines. We propose to analyze domain name system (DNS) traffic, such as Non-Existent Domain (NXDomain) queries, at several premier Top Level Domain (TLD) authoritative name servers to identify strongly connected cliques of malware related domains. We illustrate typical malware DNS lookup patterns when observed on a global scale and utilize this insight to engineer a system capable of detecting and accurately clustering malware domains to a particular variant or malware family without the need for obtaining a malware sample. Finally, the experimental results of our system will provide a unique perspective on the current state of globally distributed malware, particularly the ones that use DNS.

Summary

Primary author(s) : Dr. MOHAISEN, Aziz (Verisign Labs); THOMAS, Matthew (Verisign)

Presenter(s) : THOMAS, Matthew (Verisign)

Track Classification : Public Workshop

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

Big data journey

Sunday, 11 May 2014 10:00 (30)

On this presentation we explore the journey NZRS took to deploy and use a Big Data cluster using Hadoop.

From assembling servers, to racking, deploying software, developing UDFs and running jobs on the cluster, we go over the many alternatives of Hadoop for data analysis, and how it can be used for DNS analysis in particular.

Summary

- Building a cluster
- Running Hadoop
- Developing for Hadoop
- Data formats
- Lessons learned
- Future Work

Primary author(s) : Mr. CASTRO, Sebastian (.nz Registry Services)

Presenter(s) : Mr. CASTRO, Sebastian (.nz Registry Services)

Track Classification : Public Workshop

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

dnstap: introduction and status update

Sunday, 11 May 2014 09:30 (30)

dnstap is a flexible, structured binary log format for DNS software. This presentation will introduce the core concepts and data model and summarize recent progress in implementing dnstap support in existing DNS software.

dnstap's motivating use case is to enable an advanced form of forgery resistant passive DNS replication that can perform bailiwick verification of data received from DNS authority servers without an expensive, stateful post-processing step. This can only be done by exporting internal state from the recursive DNS server as the information that can be obtained from external packet capture is insufficient for this purpose.

However, a generic mechanism that supports the passive DNS replication use case ought to be able to support other interesting use cases. For instance, command-line tools like 'dig', 'drill', and 'kdig' produce output in similar but not identical text formats reminiscent of the DNS master file format, while various DNS "looking glass" implementations render DNS data in HTML or JSON. A unified interchange format for representing DNS transactions could substantially improve the interoperability and usability of these tools.

Summary

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

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

Track Classification : Public Workshop

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

Portable DNS Analysis

Saturday, 10 May 2014 17:20 (30)

Analyzing a DNS deployment is a complex challenge. There are several roles of DNS service, of which a single server may play multiple. Additionally, there are various vantage points from which an address might be queried, and each might result in a different response, or none at all. Finally, there are multiple query options and diverse ways handling the responses that result. There are many tools and methodologies for analyzing DNS deployments, but there is no standard, transparent way to describe the analysis or the results. We present a mechanism and framework for "portable DNS analysis" and describe its advantages for improving DNS analysis, including tool interoperability, facilitated remote analysis, and versatility.

Summary

Primary author(s) : Dr. DECCIO, Casey (Verisign Labs)

Presenter(s) : Dr. DECCIO, Casey (Verisign Labs)

Track Classification : Public Workshop

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

Analysis of DITL root data and comparison with full-resolver's data.

Sunday, 11 May 2014 16:30 (30)

The past analysis reported numbers of queries sent from each address to root DNS servers. There are 30,000 IP addresses which send over 100,000 queries in 48 hours. 100,000 queries per 48 hours seem to be too much. However, a full-resolver managed appropriately sent 110,000 queries in 48 hours at 2012 DITL timing. It served 180 queries per second from thousands of clients. The author replayed the 48 hours client queries to BIND 9 and Unbound full-resolvers, and compared number of queries to root DNS servers and other authorititative DNS servers.

Summary

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

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

Track Classification : Public Workshop

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

Performance impact of contained and virtualised environments in Authoritative DNS Servers

Sunday, 11 May 2014 15:10 (30)

Operational flexibility and deployment are increasingly managed through VMs or similar environments. In the past it has been reported that certain VM environments have a very negative impact in DNS server performance.

Here, we present the results of QPS performance of several current authoritative DNS servers running in traditional and contained or virtualised environments to evaluate their relative merits and tradeoffs in real operational use.

Summary

Primary author(s) : Mr. DAMAS, Joao (Dyn Inc); Mr. DAVE, Knight (Dyn Inc)

Presenter(s) : Mr. DAMAS, Joao (Dyn Inc); Mr. DAVE, Knight (Dyn Inc)

Track Classification : Public Workshop

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

DNS Server and DNSSEC support in Windows Server 2012 R2

Saturday, 10 May 2014 16:50 (30)

DNS Server in Windows has been enhanced significantly through recent releases of Windows Server. One of the main areas of capability augmentation of Windows DNS has been in the area of DNSSEC. This session will mainly focus on acclimatizing the user with DNSSEC capabilities in Windows DNS Server. It will demonstrate how to setup DNSSEC in Windows DNS server, online zone signing support and will provide the audience an insight into the DNSSEC validation process in Windows DNS Server. The session will also talk about Windows DNS server performance and scalability in a DNSSEC and non-DNSSEC deployment. The session will also talk about other capabilities of Windows DNS server in a file based deployment as well as Active directory based deployment.

Audience takeaways:

- DNSSEC support in Windows DNS server and how to deploy DNSSEC with Windows DNS server
- Performance and scale capabilities of Windows DNS server

Summary

Primary author(s) : Mr. ASHUTOSH, Kumar (Microsoft)

Presenter(s) : CATES, David (Microsoft)

Track Classification : Public Workshop

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

Open Resolvers in COM/NET Resolution

Sunday, 11 May 2014 17:00 (30)

While open resolvers provide various benefits by answering DNS requests from external sources for anything, today they pose a significant threat to the stability and security of the Internet. For example, open resolvers have been recently utilized for launching amplification attacks, calling for initiating a systematic study on their population, use, and distribution, and raising the awareness on their potential roles. For example, the open resolver project (http://openresolverproject.org/) reported 32 million open resolvers, 28 million of which pose a significant threat, as of October 2013.

In this presentation, we will report on an independent study of open resolvers and their usage. Beside verifying the numbers provided by the open resolver project, we go further in understanding those resolvers. To highlight their usage, we identify open resolvers in the com/net authoritative DNS resolution, and try to answer the following questions:
• What is the intersection between the open resolvers in the wild and sources of DNS requests seen in the com/net resolution?
• How persistent are the IP addresses of open resolvers in the com/net resolution over time?
• What is the correlation between the volume of DNS requests generated by open resolvers in the com/net resolution and their activity in the open resolvers ecosystem?

In this study, we received 32,040,586 responses from 31,424,854 unique IP addresses that used 277,048 forwarders. In comparison with the open resolver project statistics, and for the same time period (Oct 28, 2013 - Nov 3, 2013), our survey matched (number-wise) 98.7% of the responses and 99.03% of the unique IP addresses used by open resolvers.

We found that the daily intersection between open resolvers (forwarders) and sources of requests in the com/net resolution for the same day is more than 73% at any point in time over the time of the scan (of 7 days). Furthermore, over the time of the experiment, we found that only 91.9% of the total number of forwarders show up in the com/net resolution, with a non-trivial percent not showing up (8.1%). The daily pairwise intersection of forwarders (across different days) is shown to range from 87% to 95%, suggesting a level of dynamics and churn in the open resolvers population.

Second, we found that the number of open resolvers in the com/net resolution is persistent over time, with daily intersection ranging from 73% to 82%, and an average intersection (over 7 days) of about 76%. Third, we give each forwarder two scores: a popularity in the open resolvers ecosystem (unique IP addresses in the open resolver survey above), and a popularity score in the com/net resolution system (the number of queries issued by each forwarder). Interesting, we found that both scores are weakly and positively correlated (0.29).
Our presentation will also highlight other characteristics of open resolvers, such as geographical distribution, and persistence characterization over a longer period of time between consecutive scans (~6 months), along with implications.

Summary

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

Co-author(s) : Dr. MOHAISEN, Aziz (Verisign Labs)

Presenter(s) : WESSELS, Duane (Verisign)

Track Classification : Public Workshop

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

Zeroing in on Zero Days

Saturday, 10 May 2014 15:30 (30)

The presentation will cover findings from a Terabyte of anonymized DNS data collected every day from around the world. We'll present data and analysis techniques and discuss how we're automating the cycle of identifying and validating behaviors such as the ones described below to zero in quickly on zero days and minimize their damage.
- Appearances of new "purpose built" domains registered exclusively for amplification attacks
- A new trend of a small set of domains that go from zero (or very small) traffic and then spike to millions or 10s of millions of queries per day over a couple of days, using millions of unique random subdomains.

We'll also discuss DNS amplification attack activity at a macro and micro level.

Summary

Primary author(s) : O' LEARY, Paul (Nominum)

Presenter(s) : VAN NICE, Bruce (Nominum)

Track Classification : Public Workshop

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

Anycast on a shoe string

Saturday, 10 May 2014 16:20 (30)

Over 6 months I built out a distributed DNS service around the world consisting of 11 nodes, whilst at the same time trying to keep it under the radar of the wife - costing less than $1000/yr.

I'll talk about how I built it, what tools I used (RethinkDB, Beanstalkd, CollectD, Python etc), the problems I faced, details I learnt about how other "budget" anycast services are built and the fun I had along the way.

Summary

Primary author(s) : Mr. MORRIS, Nat (Esgob Ltd)

Presenter(s) : Mr. MORRIS, Nat (Esgob Ltd)

Track Classification : Public Workshop

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

DNSSEC Deployment in .CN

Sunday, 11 May 2014 11:30 (30)

I will introduce DNSSEC deployment in .CN in my talk, it mainly include the preparations, deployment, monitoring and observations. In the end, I will analyze a small DDoS attack occurred in .CN recently, and point out the challenges which .CN will be faced in the future.

Summary

Primary author(s) : Mr. ZHAO, Qi (CNNIC)

Presenter(s) : Mr. ZHAO, Qi (CNNIC)

Track Classification : Public Workshop

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

DNSSEC Audit Framework

Sunday, 11 May 2014 12:00 (30)

A DNSSEC audit is the process of structural examination of a DNSSEC infrastructure. DNSSEC adoption is increasing and becomes more and more a system we rely on. As the protocol becomes more critical, the level of assurance of the system and its evaluation also becomes more important.

NLnet Labs in collaboration with SWITCH created a framework that assists auditors in performing a DNSSEC audit. The framework provides a scope and a methodology, and at the same time functions as the review checklist for the audit.

Summary

Primary author(s) : MEKKING, Matthijs (NLnet Labs)

Presenter(s) : MEKKING, Matthijs (NLnet Labs)

Track Classification : Public Workshop

Contribution ID : **30** Type : **not specified**

Large scale regular expression recognition on the DITL data-set by using similarity search

Saturday, 10 May 2014 11:45 (20)

The day in the life (DITL) data-set is collected to study and improve the integrity of the root server system. Among the different properties recorded in the data-set, we focus on second level domain (SLD) strings. In this study, we introduce a method that automatically infers regular expressions from over-represented SLD strings. At first, we identify random strings and remove them from the data pipeline. Then, we find common string seeds that guide the elucidation process. Finally, we perform similarity search on strings that do not exceed a certain level of entropy level to generate a weight matrix that is then converted into regular expressions and their corresponding visualizations. Similarity search is a very expensive operation, but we manage to achieve fast results by using the simMachines R-01 similarity engine. The method may be used to preemptively discover security or performance issues in the infrastructure. During the talk, we will show a sample of collected regular expressions so that the community may identify familiar and unfamiliar SLD patterns.

Summary

Primary author(s) : Dr. MULLER-MOLINA, Arnoldo (simMachines)

Presenter(s) : Dr. MULLER-MOLINA, Arnoldo (simMachines)

Session Classification : Members-Only Session

Track Classification : Members-only

Contribution ID : **31** Type : **not specified**

OARC President Report

Saturday, 10 May 2014 12:30 (30)

Summary

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

Track Classification : Public Workshop

Contribution ID : **32** Type : **not specified**

OARC Infrastructure Update

Sunday, 11 May 2014 13:40 (30)

Summary

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

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

Track Classification : Public Workshop

Contribution ID : **33** Type : **not specified**

PGP Signing Session

Sunday, 11 May 2014 13:15 (25)

Summary

Presenter(s) : Mr. CASTRO, Sebastian (.nz Registry Services)

Track Classification : Public Workshop

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

OARC Chairman Introduction

Saturday, 10 May 2014 10:20 (20)

Summary

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

Session Classification : OARC EGM

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

Slaving the root - Warren Kumari

Sunday, 11 May 2014 16:00 (5)

Summary

Primary author(s) : Mr. KUMARI, Warren (Google)

Presenter(s) : Mr. KUMARI, Warren (Google)

Session Classification : Lightning Talks

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

Lightning Talk - Paul Vixie

Sunday, 11 May 2014 16:05 (5)

Summary

Presenter(s) : Mr. VIXIE, Paul

Session Classification : Lightning Talks

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

Lightning Talk - Matt Pounsett

Sunday, 11 May 2014 16:10 (5)

Summary

Presenter(s) : Mr. POUNSETT, Matt

Session Classification : Lightning Talks

Track Classification : Lightning talks

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

Lightning Talk - Mehmet Akcin

Sunday, 11 May 2014 16:15 (5)

Summary

Presenter(s) : Mr. AKCIN, Mehmet (Microsoft)

Session Classification : Lightning Talks

Track Classification : Lightning talks

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

Zonemaster

Sunday, 11 May 2014 16:20 (5)

Summary

Presenter(s) : Mr. WALLSTROM, Patrik

Session Classification : Lightning Talks

Track Classification : Lightning talks

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

Standarized DNS measurement

Sunday, 11 May 2014 16:25 (5)

Summary

Presenter(s) : Mr. MARTIN, Jim; Mr. LIMAN, Lars-Johan

Session Classification : Lightning Talks

Track Classification : Lightning talks