


Prefix Top Lists: Gaining Insights with Prefixes from Domain-based Top Lists on DNS Deployment

Johannes Naab, Patrick Sattler, Jonas Jelten, Oliver Gasser, Georg Carle
<lastname>@net.in.tum.de

Tuesday 22nd October, 2019

ACM Internet Measurement Conference 2019, Amsterdam, Netherlands

Chair of Network Architectures and Services
Department of Informatics
Technical University of Munich



- Existing **domain based top lists** (e.g. Alexa, Majestic, Umbrella) proved to be of high value for research
- So far work on improving top lists mainly focused on providing **more stable** domain top lists
- Evaluations focus on domain based top lists, but not referenced Internet resources

Our Goals:

- Provide method to generate new top list types: rank **prefixes** and **ASes** as important Internet resources
 - by assigning weights on the domain based top list elements
 - using a comprehensive name resolution process
- Provide a metric to investigate changes in Prefix and AS top lists
- Show usefulness of new top lists when analyzing DNS resilience

Outline

Problem Statement

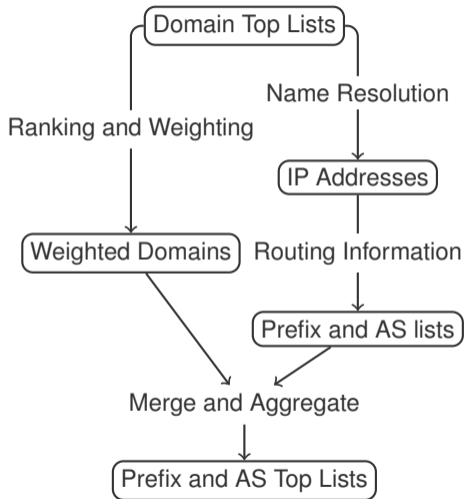
Methodology

Prefix Top List Instantiation


Prefix Top List Comparison Metric

DNS Resilience

Conclusion



Prefix Top List Instantiation



Domain Top Lists

- Alexa, Majestic, Umbrella

Ranking and Weighting


- Zipfian distribution to approximate popularity weights

$$w = \frac{1/k}{\sum_{n=1}^N 1/n}, k = \text{rank}, N = \text{number of elements}$$

- smoothing of fluctuations: 7-day aggregate

Rank	Domain	Weight	Top Rank	Bottom Rank
1	google.com	0.0703	1	1
2	youtube.com	0.0351	2	2
3	tmall.com	0.0226	3	4
4	baidu.com	0.0184	3	4
5	qq.com	0.0134	5	6
6	sohu.com	0.0120	5	8
7	facebook.com	0.0099	7	8

Top Domains for August 1, 2019 with 7-day rolling window, Alexa List



Name Resolution

- our instantiation: local vantage point
- query A and AAAA records
- against all authoritative name servers

Routing Information

- our instantiation: local BGP view
- map resolved addresses to prefix and origin AS

Merge and Aggregate

- calculate prefix weight from domain weights
- if DNS resolution returns multiple IP addresses: split
 - preserve invariant: $\sum(w) = 1$

Rank	Prefix	Weight
1	172.217.18.0/24	0.0178
2	172.217.16.0/24	0.0175
3	172.217.22.0/24	0.0173
4	216.58.206.0/23	0.0165
5	172.217.23.0/24	0.0164
6	140.205.64.0/18	0.0160
7	216.58.208.0/24	0.0154
8	111.160.0.0/13	0.0134
9	140.205.128.0/18	0.0116
10	216.58.204.0/23	0.0098
11	172.217.21.0/24	0.0097
12	39.156.0.0/17	0.0096
13	216.58.210.0/24	0.0094
14	220.181.32.0/19	0.0092
15	91.198.174.0/24	0.0073
16	151.101.0.0/22	0.0063
17	151.101.64.0/22	0.0062
18	151.101.192.0/22	0.0062
19	151.101.128.0/22	0.0061
20	99.84.88.0/21	0.0057

BGP Prefix Ranking for August 1, 2019 based on Alexa List.

Rank	Prefix	Weight	# Domains	# IP addr.
1	172.217.18.0/24, AS15169 – GOOGLE	0.0178	1039	35
2	172.217.16.0/24, AS15169 – GOOGLE	0.0175	1000	33
3	172.217.22.0/24, AS15169 – GOOGLE	0.0173	1041	42
4	216.58.206.0/23, AS15169 – GOOGLE	0.0165	973	35
5	172.217.23.0/24, AS15169 – GOOGLE	0.0164	775	23
6	140.205.64.0/18, AS37963 – CNNIC-ALIBABA	0.0160	6	4
7	216.58.208.0/24, AS15169 – GOOGLE	0.0154	443	14
8	111.160.0.0/13, AS4837 – CHINA169-BACKBONE	0.0134	3	4
9	140.205.128.0/18, AS37963 – CNNIC-ALIBABA	0.0116	12	12
10	216.58.204.0/23, AS15169 – GOOGLE	0.0098	547	15
11	172.217.21.0/24, AS15169 – GOOGLE	0.0097	697	22
12	39.156.0.0/17, AS9808 – CMNET-GD	0.0096	7	3
13	216.58.210.0/24, AS15169 – GOOGLE	0.0094	403	12
14	220.181.32.0/19, AS23724 – CHINANET-IDC-BJ-AP	0.0092	9	5
15	91.198.174.0/24, AS14907 – WIKIMEDIA	0.0073	12	1
16	151.101.0.0/22, AS54113 – FASTLY	0.0063	4566	192
17	151.101.64.0/22, AS54113 – FASTLY	0.0062	4475	183
18	151.101.192.0/22, AS54113 – FASTLY	0.0062	2157	187
19	151.101.128.0/22, AS54113 – FASTLY	0.0061	2136	182
20	99.84.88.0/21, AS16509 – AMAZON-02	0.0057	11 988	168

BGP Prefix Ranking for August 1, 2019 based on Alexa List.

AS Popularity Rank and Weight

Rank AS	Rank Prefix	Object	Weight	# Domains	# IP addr.
	1	172.217.18.0/24	0.0178	1039	35
	2	172.217.16.0/24	0.0175	1000	33
	...				

IPv4 Object Ranking for August 1, 2019 based on Alexa List.

AS Popularity Rank and Weight

Rank AS	Rank Prefix	Object	Weight	# Domains	# IP addr.
1		AS15169 – GOOGLE	0.1454	103 264	15 278
	1	172.217.18.0/24	0.0178	1039	35
	2	172.217.16.0/24	0.0175	1000	33
	...				

IPv4 Object Ranking for August 1, 2019 based on Alexa List.

AS Popularity Rank and Weight

Rank AS	Rank Prefix	Object	Weight	# Domains	# IP addr.
1		AS15169 – GOOGLE	0.1454	103 264	15 278
	1	172.217.18.0/24	0.0178	1039	35
	2	172.217.16.0/24	0.0175	1000	33
	...				
2		AS13335 – CLOUDFLARE	0.1049	310 574	91 869
	28	23.227.38.0/23	0.0045	42 809	13
3		AS16509 – AMAZON-02	0.0651	88 373	70 888
	20	99.84.88.0/21	0.0057	11 988	168
4		AS37963 – CNNIC-ALIBABA	0.0478	7266	6733
	6	140.205.64.0/18	0.0160	6	4
5		AS54113 – FASTLY	0.0284	16 752	887
	16	151.101.0.0/22	0.0063	4566	192
	...				
14		AS20940 – AKAMAI-ASN1	0.0106	3201	2807
	80	23.38.48.0/20	0.0016	203	104
15		AS32934 – FACEBOOK	0.0102	48	38
	24	157.240.20.0/24	0.0051	36	5
16		AS26496 – GO-DADDY	0.0082	116 590	24 431
	112	184.168.128.0/22	0.0011	11 538	2
17		AS14907 – WIKIMEDIA	0.0073	26	4
	15	91.198.174.0/24	0.0073	12	1

IPv4 Object Ranking for August 1, 2019 based on Alexa List. Full data on <https://prefixtoplists.net.in.tum.de/>.

Different lists can be compared by the sum of the weight differences.

Example: tum.de


- 2019-07-31: rank = 12 556, $w = 5.32 \times 10^{-6}$
- 2019-08-01: rank = 13 593, $w = 4.92 \times 10^{-6}$
- $\Delta = 0.40 \times 10^{-6}$

Different lists can be compared by the sum of the weight differences.

Example: tum.de


- 2019-07-31: rank = 12 556, $w = 5.32 \times 10^{-6}$
- 2019-08-01: rank = 13 593, $w = 4.92 \times 10^{-6}$
- $\Delta = 0.40 \times 10^{-6}$

Difference between two lists: $\text{change} = \sum_{i \in \text{elements}} \text{abs}(\Delta_i)$

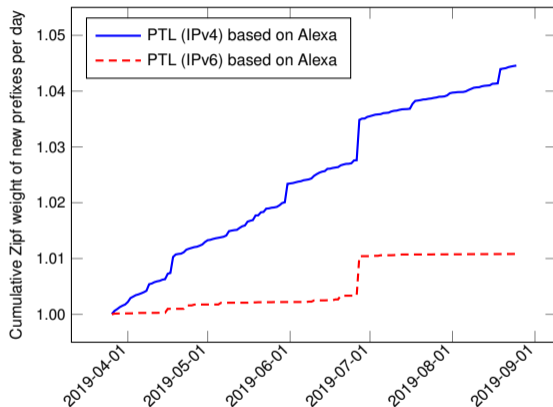


Daily weight changes for Alexa Domain and Prefix Top List (PTL) with 7 day aggregate.

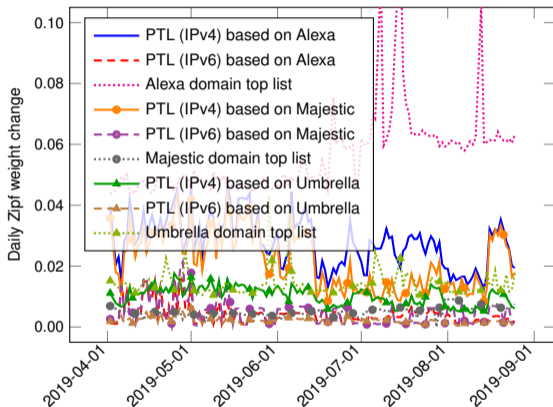
Stability of Ranking



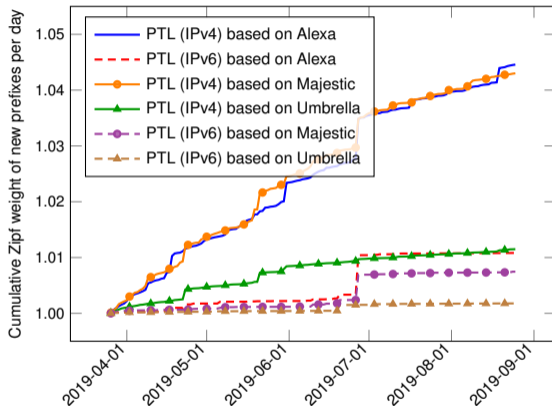
Daily weight changes for Alexa Domain and Prefix Top List (PTL) with 7 day aggregate.



Cumulative weight of new prefixes. The jump on June 27, 2019 is caused by new prefixes for *Wikipedia*.



Daily weight changes for Alexa Domain and Prefix Top List (PTL) with 7 day aggregate.



Cumulative weight of new prefixes. The jump on June 27, 2019 is caused by new prefixes for *Wikipedia*.

Check for RFC 2182 compliance, i.e. name servers replicas in distinct /24, resp. /48 prefixes, using our name resolution results.

Allman analyzed this last year at IMC and found $\approx 12\%$ to not fulfilling this requirement.

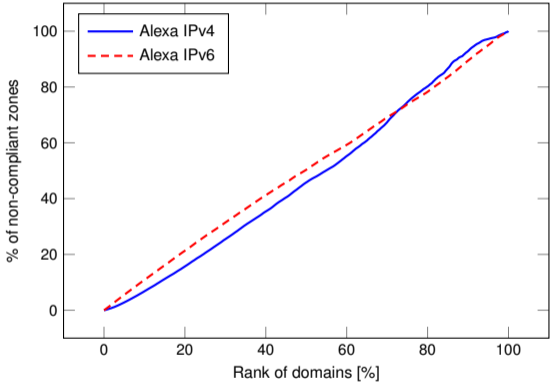
Check for RFC 2182 compliance, i.e. name servers replicas in distinct /24, resp. /48 prefixes, using our name resolution results.

Allman analyzed this last year at IMC and found $\approx 12\%$ to not fulfilling this requirement.


Using our active measurements we find similar results:

	Alexa IPv4	Alexa IPv6
Domains	2.9M	1.4M
Non-compliant	411k	610k
Share	14.26%	43.54%

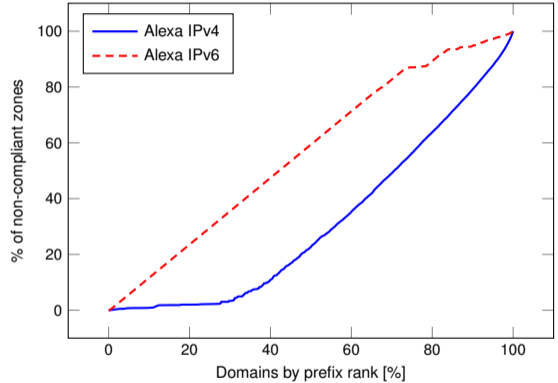
Question: Are popular domains more resilient?



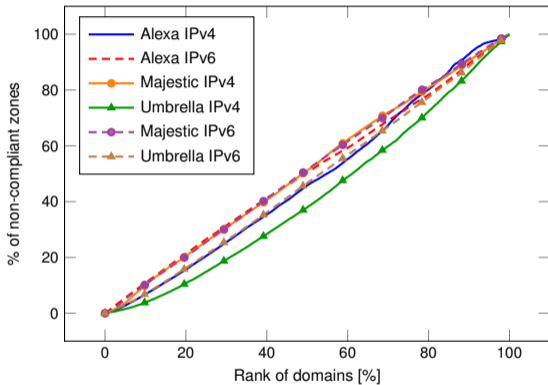
Non-compliant zones ordered by their rank, x-axis normalized to number of resolved domains.




Non-compliant zones ordered by their rank, x-axis normalized to number of resolved domains.



Non-compliant zones ranked by the prefix they are in, x-axis normalized to the number of resolved domains.



Non-compliant zones ordered by their rank, x-axis normalized to number of resolved domains.

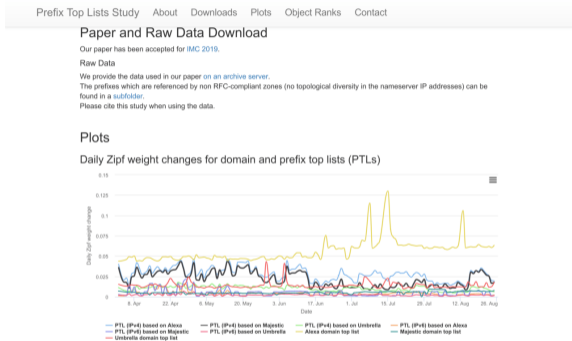


Non-compliant zones ranked by the prefix they are in, x-axis normalized to the number of resolved domains.

- Provide a method to create Prefix Top Lists from domain top lists
- Approximate popularity weights by Zipfian distribution
- Prefix Popularity Ranking and AS Popularity Ranking based on comprehensive Name Resolution
- Metric to quantify change based on weight and composition changes
- Analysis of DNS resilience via Prefix Top Lists shows centralization towards popular prefixes

Outlook: Investigate additional instantiations

- Global list: multiple vantage points
- Investigate impact of DNS load balancing



<https://prefixtoplists.net.in.tum.de/>