A Farsight Security® Report

FARSIGHT SECURITY

Joe St Sauver, Ph.D.
Distinguished Scientist, Farsight Security, Inc.

DNS Network Traffic Volumes During the 2020 Pandemic

New Farsight Security Study Provides Results for over 300 Domains (Including Possible DDoS Event Traffic)

Executive Summary

Context: The COVID-19 coronavirus pandemic of 2020 has changed how users interact with the world and particularly with the Internet. Employees may now be working from home (or may have been laid off as a result of government-mandated business closures). Colleges and universities have largely shuttered their campuses and gone online. Online shopping has supplanted visits to brick-and-mortar stores. Business and leisure travel has largely ceased. Binge-watching streaming television has replaced going out for dinner, drinks, and a traditional movie. The world we inhabit today is NOT the same world we inhabited at the end of 2019.

Study Focus: This study, which took place from March-April 2020, reviews Farsight Security's observed traffic volumes for passive DNS cache miss traffic for 316 online sites selected from five broad categories: (1) news, (2) travel and transportation, (3) retail, (4) streaming video and (5) colleges and universities.

How We Did It: We graphed changes in DNS cache miss traffic using "volume over time" code from an earlier Farsight Security blog post. That code leverages Farsight-archived daily MTBL-format data files, each of which has per-day counts for each unique (RRset, RRtype, Bailiwick, Rdata) combination seen during the file's 24-hour period. Graphs of the raw data with an overlaid smoothed 7 day moving average were produced for each of the 316 studied sites.

Among the Key Findings

- A "step up" pattern in traffic typically reflecting a 4x-to-7x increase in DNS cache miss traffic levels was seen across most or all verticals. This change took place, often abruptly, between mid-to-late-March and early-to-mid-April.
- While most of the studied sites exhibited this “step up” traffic pattern, there was variation among the studied sites (for example, higher education tends to exhibit a gradual increase, and that increase then drops again, producing a hill rather than a plateau). The report provides individual graphs for each of the sites.
- Some sites experienced "spikes" in volume, spikes that were so large that the spikes caused most of the "normal variation" in traffic volume to "wash out" due to the dominance of the spike or spikes.
- Farsight believes those spikes represent denial of service (DDoS) attack traffic reflexively targeting some unrelated third-party site or sites.
    - At least two distinct reflective distributed denial of service attack patterns apparently took place among the studied sites: One type that appears to be purely associated with abusive DNS SOA ("Start of Authority") queries, and
    - A second type that melds abusive DNS SOA queries with abusive DNS TXT queries for wildcarded SPF redirect records

This report does not try to "attribute" or "apportion" the change in traffic levels to all the potential sources mentioned below. Instead, it simply reports what we see as a macroscopic phenomenon, and invite authoritative reports from other data sources (such as those with flow level data for particular sites).

Changes in observed traffic levels may be the result of many potential interacting and amplifying (or modulating) factors:

- Users may simply be more active online (perhaps searching diverse sources for the latest news about coronavirus or trying to shop online).
- Users may have moved their Internet visits from old (poorly instrumented) locations to new (better-instrumented) ones.
- Internet sites may be changing their service or content, perhaps moving to a content distribution network for increased capacity, or increasing their online advertising in an attempt to backfill revenue lost from other sources. Authoritative nameserver operators may have made technical changes to TTL (time to live) values, or recursive resolver operators may have built out instrumented resolver farms to better handle the new load they're facing from their users.
- Some changes in DNS traffic may also be due to denial of service attack traffic or hacking/cracking attempts (as in the case of some “spikes” in traffic we observed.
- At the same time, DNS cache-related effects may mask or modulate changes in observed traffic volumes (since a single cache miss seen above a large shared recursive resolver may actually represent hundreds, thousands, or even millions of downstream user cache hits answered directly out of the resolver's cached data).

Table of Contents

Section I. Overview

1) Why Look at Domain Name Counts?

Farsight's passive DNS database, DNSDB, tracks, catalogs and indexes the unique **relationships** present in DNS resource records. That fits well with the way most users employ our data -- they tend to be interested in "what points to what" or perhaps the IP address hosting history for a domain. However, what's sometimes overlooked is the fact that Farsight also sees and reports the **count** for the cache miss traffic reported by its sensors.

Those counts are normally aggregated into a single value for each unique (RRname, RRtype, Bailiwick, Rdata) value. Those counts represent the aggregate number of observations seen across the full first-to-last seen period for that combination, a period which can literally be up to a decade in some cases. As an example, here's the count for one of Farsight's own domains, showing that it represents 2,133,799 observations going from July 2013 right up to the date the query was run:

```
$ dnsdbq -r farsightsecurity.com -S -k count
;; record times: 2013-07-17 21:26:20 .. 2020-05-08 22:46:14
;; count: 2133799; bailiwick: farsightsecurity.com.
farsightsecurity.com.  NS  ns5.dnsmadeeasy.com.
farsightsecurity.com.  NS  ns6.dnsmadeeasy.com.
farsightsecurity.com.  NS  ns7.dnsmadeeasy.com.
[etc]
```

Sometimes, though, you may want **per day counts** for each unique (RRname, RRtype, Bailiwick, Rdata) combination.

Some deaggregated data is available via DNSDB's "gravel" feature,[1] but that format may include data in time "chunks" of different size (e.g., yearly values, monthly values, daily values, hourly values, etc.), and you may still need to do some data reduction on your own, depending on what you're after.[2]

Fortunately, we can also get direct daily counts if we have access to individual daily MTBL files -- each of those files reports the count for each unique (RRname, RRtype, Bailiwick, and Rdata) combination, but those counts have NOT been aggregated Each daily has a count that represents just the counts for **that day's worth of traffic.**

We're going to leverage our archive of individual daily MTBL files to report on the volume of DNS cache miss traffic we've seen for selected domains during the current coronavirus pandemic.[3]

We were inspired to undertake this project in part because of a report in the New York Times.[4] That NY Times article reported changes in the number of "sessions" for both selected web sites and selected mobile apps during the pandemic, and illustrated their findings with a small selection of graphs.

We wanted to look at a different measure, and look more broadly. It wouldn't be sufficient to simply look at the macroscopic size of our deduplicated datastreams -- because Farsight simply increments a counter when it sees an additional "hit" for a previously seen (RRname, RRtype, Bailiwick, Rdata)-tuple, data volumes for common names could double, increase by a factor of ten, or increase by a factor of a 100 or more with no ripple in the macroscopic size of the datastream.

---

[1] "Crushing Monolithic Data Results ("Rock") Into "Gravel": dnsbq New -g Volume-Across-Time Option," https://www.farsightsecurity.com/blog/txt-record/gravel-20190927/

[2] For example, if you want to know the total number of NS record "hits" per day across all the various Rdata values that may have been seen, you'd need to roll those up yourself to get the number you're actually after.

[3] For an example of how to process MTBL files to extract data of interest, see "Finding Top FQDNs Per Day in DNSDB Export MTBL Files (Part One of a Three Part Series)" https://www.farsightsecurity.com/blog/txt-record/TopFQDNs-20190322/

[4] "The Virus Changed the Way We Internet," https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html

To get a sense of traffic increase or decrease, **you need to look at the counts associated with each relevant record**. **For a series of different domains, we wanted to see if the aggregate volume of cache miss traffic -- as seen in the counts for all hosts under each of a series of domains -- changed during the COVID-19 pandemic.**[5]

We suspect that we should first backup and explain what we mean by "cache miss traffic."

Remember that most Internet users go to a relatively small number of highly popular sites -- Google, Facebook, Twitter, Amazon, Netflix, etc. Given that behavior, it doesn't make much sense for an ISP's recursive resolvers (the ISP's "name servers") to repeatedly ask the Internet to answer the same question "from scratch" a gazillion times a second. Everything will be much more efficient if the large ISP's name servers just "cache" ("remember") the answers they've recently seen for popular queries, and then answer those queries (when users repeatedly ask them again) from the locally-cached values. (These answers-from-the-cached-values are called "cache hits.")

If the user's query is one for a name that **hasn't** been seen and cached recently, the recursive resolver must then chase down the information the user requires. That's called a "cache miss." Cache miss traffic is what Farsight's sensors watch at hundreds of sites all around the Internet, and that's what we're going to report on in this report. See the following diagram:

Figure 1. Farsight Collects Cache Miss Traffic Above Large Recursive Resolvers

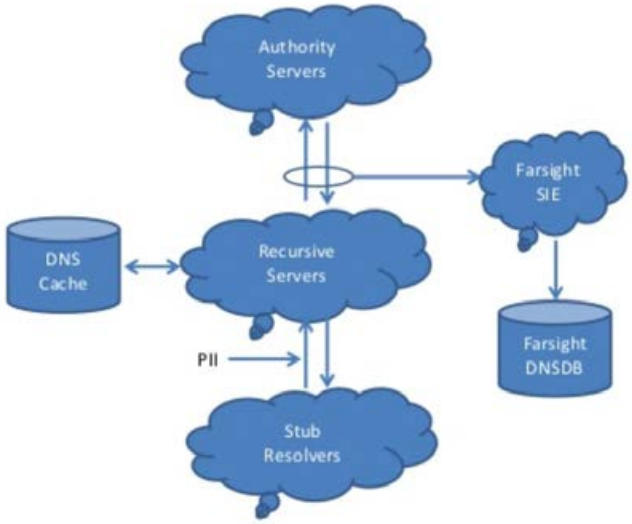

While our sensors don't (and can't) see ALL cache miss traffic worldwide, we see ENOUGH cache miss traffic to get a pretty good idea about what's going on when it comes to DNS traffic levels.

Collecting cache miss traffic above large recursive resolvers ensures that Farsight does NOT see the personally identifiable information that might otherwise be exposed between the stub resolver and the recursive resolver.

[5] For a timeline of some relevant current events during this time period, please see Appendix I.

2) We May Be Seeing Actual Changes in User Activity... Or Some Other Things

Many different factors can influence the aggregate volume of cache miss traffic Farsight sees for a domain over time, including:

- **End User Online Behavior May Actually Have Changed:** Users may visit different sites than they normally do, or may visit the same sites more (or less) often than normal. **This is what most folks are typically interested in, what they're hoping we are largely reporting here, although other factors most likely also influence the reported counts, including....**

- **Impacts of Potential TTL (Times To Live) Changes:** When a caching DNS resolver gets an answer for a DNS query, the domain owner will suggest how long the caching resolver should remember and rely on that answer. This time is known as the "time to live" or "TTL", and is measured in seconds. Once an answer has been cached, subsequent queries for that same name and record type will be answered using that cached value until the TTL counts down to zero and the cached response is discarded.

  TTLs often are reduced when a domain owner is planning to make changes to their DNS, since a short TTL will ensure that any changes the domain owner makes will rapidly be propagated Internet-wide. On the other hand, TTLs may be increased when name servers are heavily loaded and there's a desire to reduce load on the existing servers.

  Farsight does not currently collect and report effective TTLs for individual resource records, so we cannot precisely identify changes to TTLs that may have occurred -- *if* any have indeed changed during the study period. That said, we don't expect that all (or even a material fraction of) the sites we studied will have made TTL changes during the study period.

- **Sensor Changes:** Farsight collected passive DNS data from sites all around the world. As a matter of company policy, Farsight does not disclose the identities of its sensor operators, nor the location of individual sensors.

  That said, changes to Farsight's overall sensor footprint do happen from time -- for example, new sensors may come online, and/or one or more established sensors may go offline (whether temporarily for maintenance, or permanently). These sort of changes may potentially result in changes to traffic levels. Please note, however those changes in volume are normally incremental (rather than orders of magnitude in size). We don't believe that sensor changes explain the difference in cache miss traffic volume we observed during the study period.

- **Destination Site Changes May Have Taken Place:** For instance, Internet sites may be changing their service or content, perhaps moving to a content distribution network for increased capacity, or increasing their online advertising in an attempt to backfill revenue lost from other sources. Sites may also have reduced content due to staff layoffs or other factors (for example, a sports site might not have current games to cover due to sport seasons having been canceled).

- **Changes to User Locations:** Even if our sensor footprint doesn't change, users may change **where they're working from,** potentially going from a "poorly instrumented location" (where we traditionally haven't seen their cache miss traffic) to a new well-instrumented location (where we do).

- **We May Not Be Seeing All Of The Impact of End User Behavior Changes Due to Caching:** Finally, caching may also mask the true magnitude of volumetric changes. A single cache miss seen by Farsight's sensor code may represent one user's interest in a site, or that of a hundred, a thousand, or a million users. We just don't have the ability to tell given that most queries at the recursive resolver will be answered from the recursive resolver's cached traffic.

- **Some changes in DNS traffic** may also be due to denial of service attack traffic or hacking/cracking attempts (as in the case of some “spikes” in traffic we observed.

Even given the multiple potential challenges just mentioned, we still thought it would be worth looking to see if (and how) DNS cache miss traffic levels have changed during the pandemic. We'll therefore share the graphs in this report on an "as-is" basis, subject to all the limitations described above.

3) The Categories of Sites We Reviewed

We decided to look at four main categories of online sites for this report:

- **News Sites.** Many users have an increased appetite for news, searching for pandemic-related information, or information about when the country will be able to reopen.

  We've included both leading television and print media sources, and popular "partisan" sites with either a distinct liberal ("left wing") or distinct conservative ("right wing") perspective. The initial New York Times report[6] that served to inspire this work saw a difference in volume change by type of news site, and we became curious to see if partisanship would impact what we saw as well (for what it may be worth, our data shows an uptick in interest across news sites with only rare exceptions).

- **Travel and transportation-related sites**. This category includes airlines, car rental companies, cruise ship companies, hotels, railroad lines and trucking companies. We expected lower traffic to these sites since business and leisure travel has been at a standstill, but many may be working to rebook or refund pre-existing reservations, or may be working to arrange the logistics for emergency supplies such as personal protective equipment, ventilators, medications, etc.

  Because airline flight volume has dropped 95% from pre-COVID volumes,[7] we expected to see particularly broad and deep drops in airline-related DNS cache miss traffic. Because that was not the case for an initial set of sites we investigated, and to get a sense of how a broader set of airline domains were impacted, we also checked many other international carriers and even some cargo airlines. We developed candidate sites for that expanded group based on a number of sites[8] highlighting leading airlines abroad, and eventually graphed a total of 73 different airlines as shown in the body of the report.

- **Retail sites**: We expected these sites would be impacted by closure orders and a general move to online ordering for delivery. Our retail sites category includes both large box retailers, grocery chains, online retailers, pharmacies, some of the country's largest fast food chains, etc.

- **Higher education sites**: Virtually all colleges and universities have closed their physical campuses and moved instruction and research exclusively online. The sites we're showing here includes a selection of leading national universities (public and private), the Ivy League, an assortment of liberal arts colleges, as well as a selection of leading international universities.

  We included this range of institutions because we know that how things work at a large state school is often dramatically different than how things work at a small liberal arts college or an international university.

[6] "The Virus Changed the Way We Internet,"
https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html

[7] "Major US airlines to require facial coverings on flights in May,"
https://www.cnbc.com/2020/05/01/us-airlines-american-delta-and-united-to-require-facial-coverings-on-flights.html

[8] https://en.wikipedia.org/wiki/Largest_airlines_in_the_world
https://www.iata.org/contentassets/a686ff624550453e8bf0c9b3f7f0ab26/wats-2019-mediakit.pdf
https://www.statista.com/statistics/697370/largest-airlines-in-europe-by-passengers/
among others.

4) Methodology -- What EXACTLY Did We Look at For Each Site?

When reviewing traffic for these sites, we looked at the DNS cache miss traffic Farsight sensors saw for ALL hostnames under a given delegation point (e.g., *.example.com, **NOT**, for example, just a specific hostname such as www.example.com).

We also looked at ALL resource record types[9] (excluding only DNSSEC-related record types). This means the counts you see include IPv4 "A" records, IPv6 "AAAA" records, CNAMEs, NS records, MX records, TXT records, SOA records, etc.

For the purpose of this report, we limited the bailiwick[10] to just the delegation point level (we did not look at data from the TLD servers or the root).

We processed the data for the graphs in Section II using the "VoT" code previously described in a Farsight Security blog article from 2019.[11]

A typical run looked like:

```
$ vot --fqdn \*.nscorp.com --rectype ANY --bailiwick nscorp.com --ma 7 --plot
--plotma --device postscript --start 20200301 --stop 20200430
```

Each of the graphs in Section II has two lines:

- One line, dashed, represents the actual observed count values per day
- The other solid line is a smoothed seven day moving average (that smoothing reduces the impact of day-to-day fluctuations on the graphed line).

At the risk of stating the obvious, most of the graphs show a consistent pattern: initially levels are relatively low, then levels quickly increase to a new higher value that's often ~5x the former level.

For example, looking at the actual numerical data for *.forbes.com, that data averaged 61,342.7 through 2020-03-31, but 335,217.9 thereafter, an increase of ~5.5x. (We'll look at the *.forbes.com example in more depth in the next section.)

Not all sites increased by the same ratio. For example:

- *.foxnews.com went from 229,501.3 through 2020-03-31, up to 855,972.4, an increase of 3.7x
- *.newegg.com went from 42,698.5 through 2020-03-31, up to 261,860.8, an increase of 6.1x
- *.chewy.com went from 15,361.5 through 2020-03-31, up to 132,973.5, an increase of 8.6x
- *.bnsf.com went from 5762.3 through 2020-03-31, up to 81326.5, an increase of 14.1x
- *.hertz.com went from 24,598.7 through 2020-03-31, up to 702,860.8, an increase of 28.6x

[9] "A Quick Overview of the Top Seven DNS Record Types,"
https://www.farsightsecurity.com/blog/txt-record/dnsrecords-20171201/

[10] "What is a Bailiwick?"
https://www.farsightsecurity.com/blog/txt-record/what-is-a-bailiwick-20170321/

[11] "Volume-Over-Time Data From DNSDB Export MTBL Files (Part Two of Three-Part Series)",
https://www.farsightsecurity.com/blog/txt-record/volumeovertime-20190401/

5) Digging In On a Specific Example

Consider the change in volume shown for *.forbes.com:

Figure 1. *.forbes.com Daily Traffic Volume Over Time

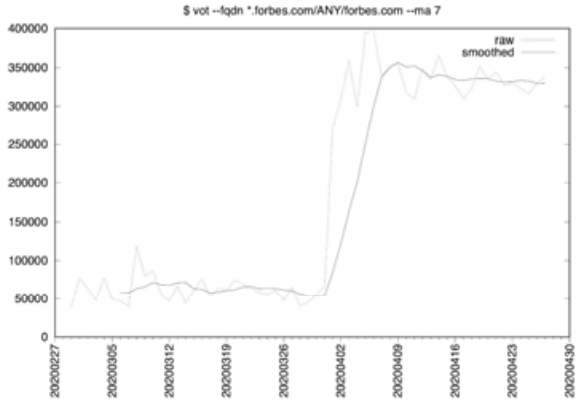

Figure 1 exhibits a fairly flat initial region, then abruptly "steps up," going from a count of about 61,342.7 queries/day (prior to April 1) to 5.5x that level afterwards.

What specific fully qualified domain names (FQDN)/RRtypes driving that traffic volume change? Is it all interactive visits to www.forbes.com? Something else, such as infrastructural traffic associated with the Forbes nameservers?

Looking at a typical day from the initial "lower volume" region, we see:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200319.D.mtbl

$ dnstable_lookup -j rrset \*.forbes.com any forbes.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | wc -l
58                      <-- 58 unique RRsets found

$ dnstable_lookup -j rrset \*.forbes.com any forbes.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | sort -nr | less
28860 forbes.com. NS ["ns-217.awsdns-27.com.","ns-979.awsdns-58.net.","ns-1028.awsdns-
00.org.","ns-1637.awsdns-12.co.uk."]
10786 fast.forbes.com. CNAME ["g2.shared.global.fastly.net."]
4549 fuse.forbes.com. CNAME ["d.sni.global.fastly.net."]
3676 email.forbes.com. NS ["ns1.crdl.io.","ns2.crdl.io.","ns3.crdl.io.","ns4.crdl.io."]
3583 blogs-images.forbes.com. CNAME ["n2.shared.global.fastly.net."]
3090 forbes.com. A ["151.101.2.49","151.101.66.49","151.101.130.49","151.101.194.49"]
2750 forbes.com. SOA ["ns-1637.awsdns-12.co.uk. awsdns-hostmaster.amazon.com. 1 7200 900
1209600 86400"]
1437 geolocation.forbes.com. CNAME ["g2.shared.global.fastly.net."]
1001 www.forbes.com. CNAME ["g2.shared.global.fastly.net."]
661 aax.forbes.com. CNAME ["lb.aaxads.com."]
624 thumbor.forbes.com. CNAME ["g2.shared.global.fastly.net."]
592 images.forbes.com. CNAME ["g2.shared.global.fastly.net."]
484 related.forbes.com. CNAME ["forbes.media.net."]
384 damapi.forbes.com. NS ["ns-cloud-b1.googledomains.com.","ns-cloud-
b2.googledomains.com.","ns-cloud-b3.googledomains.com.","ns-cloud-b4.googledomains.com."]
192 www3.forbes.com. CNAME ["www3.forbes.com.edgekey.net."]
160 login.forbes.com. A ["151.101.1.195","151.101.65.195"]
```

Nothing there looks particularly odd to our casual inspection.

Now what about a day from the "higher plateau" region? We see similar RRnames, but at a higher level:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200415.D.mtbl
$ dnstable_lookup -j rrset \*.forbes.com any forbes.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | wc -l
82
$ dnstable_lookup -j rrset \*.forbes.com any forbes.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | sort -nr | head -16
168445 forbes.com. NS ["ns-217.awsdns-27.com.","ns-979.awsdns-58.net.","ns-1028.awsdns-
00.org.","ns-1637.awsdns-12.co.uk."]
80985 fast.forbes.com. CNAME ["g2.shared.global.fastly.net."]
21169 fuse.forbes.com. CNAME ["d.sni.global.fastly.net."]
20898 email.forbes.com. NS ["ns1.crdl.io.","ns2.crdl.io.","ns3.crdl.io.","ns4.crdl.io."]
18400 blogs-images.forbes.com. CNAME ["n2.shared.global.fastly.net."]
12495 forbes.com. SOA ["ns-1637.awsdns-12.co.uk. awsdns-hostmaster.amazon.com. 1 7200 900
1209600 86400"]
7254 aax.forbes.com. CNAME ["lb.aaxads.com."]
5186 images.forbes.com. CNAME ["g2.shared.global.fastly.net."]
4971 forbes.com. A ["151.101.2.49","151.101.66.49","151.101.130.49","151.101.194.49"]
3769 login.forbes.com. A ["151.101.1.195","151.101.65.195"]
2966 geolocation.forbes.com. CNAME ["g2.shared.global.fastly.net."]
2453 damapi.forbes.com. NS ["ns-cloud-b1.googledomains.com.","ns-cloud-
b2.googledomains.com.","ns-cloud-
b3.googledomains.com.","ns-cloud-b4.googledomains.com."]
2067 li.forbes.com. CNAME ["4635a7195419e75d73346b624e7372f6.edgesuite.net."]
1975 nyccbpro01.forbes.com. A ["172.17.2.38"]
1415 forbesmags.forbes.com. CNAME ["go.pardot.com."]
1202 www.forbes.com. CNAME ["g2.shared.global.fastly.net."]
```

Let's pull 16 of those FQDNs and associated count values into a consolidated table for ease of comparison (the "subfigure" letter values refer to graphs that are part of Figure 3, below).

Table 1. DNS Cache Miss Query Counts For Two Comparative Dates, Selected *.forbes.com FQDNs

Subfigure	RRset	Count (2020-04-15)	Count (2020-03-19)	Ratio
a)	forbes.com. NS	168445	28860	5.83
b)	fast.forbes.com. CNAME	80985	10786	7.50
c)	fuse.forbes.com. CNAME	21169	4549	4.65
d)	email.forbes.com. NS	20898	3676	5.68
e)	blogs-images.forbes.com. CNAME	18400	3583	5.13
f)	forbes.com. SOA	12495	2750	4.54
g)	forbes.com. A	4971	3090	1.60
h)	geolocation.forbes.com. CNAME	2966	1437	2.06
i)	www.forbes.com. CNAME	1202	1001	1.20
j)	aax.forbes.com. CNAME	7254	661	**10.97**
k)	thumbor.forbes.com. CNAME	970*	624	1.55
l)	images.forbes.com. CNAME	5186	592	8.76
m)	related.forbes.com. CNAME	597*	484	1.23
n)	damapi.forbes.com. NS	2453	384	6.38
o)	www3.forbes.com. CNAME	395*	192	2.05
p)	login.forbes.com. A	3769	160	**23.55**

* = manually looked up since not in the shown default "top 16" list for this date

If we focus on the "ratio" column in the preceding table, two things are apparent:

- The number of DNS cache miss queries associated with "login.forbes.com" has increased tremendously (by a factor of nearly 24x). This would be consistent with more subscriber-only content requiring authentication.

- The number of DNS cache miss queries associated with "aax.forbes.com" has also increased by a factor of nearly 11x. aax.forbes.com is a CNAME that points at lb.aaxads.com, an online advertising company, see https://aax.media

These volume changes would be consistent with Forbes creating and effectively monetizing highly in-demand content for subscribers, such as information about the coronavirus pandemic.

But is the coronavirus a true "traffic magnet" topic? Yes. Check out the following Google Trends chart (we include "trump" and "biden" as reference comparators -- obviously those two normally-highly-newsworthy individuals are "way down in the weeds" in comparison to interest in "coronavirus"):

Figure 2. Google Trends Report Showing Comparative Interest
in "coronavirus" vs. The Two Major U.S. Presidential Candidates

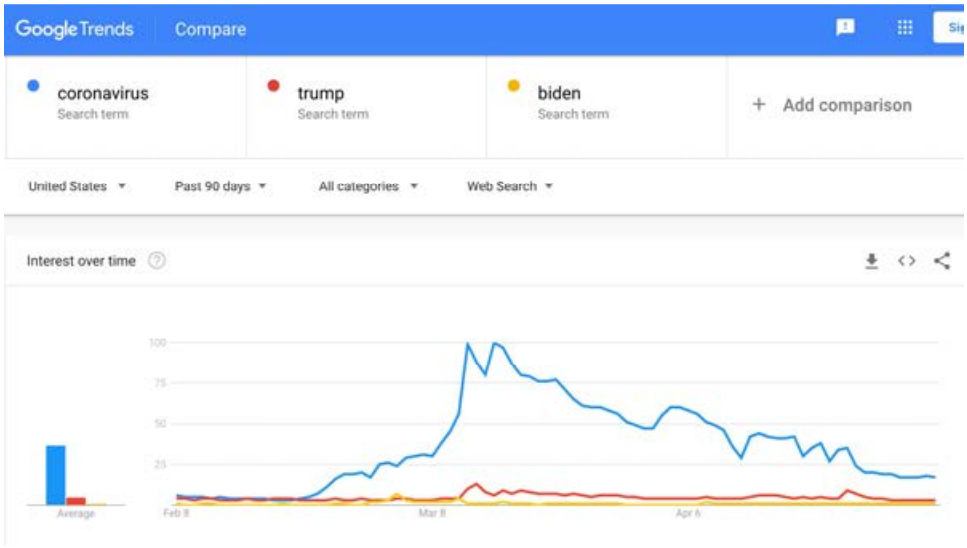

We'll now show you the subfigures that correspond to the letters from Table 1, above. These "finer-grained" graphs will serve to "disect" and "isolate" the traffic that is represented in aggregated form in Figure 1, above.

Much of the traffic looks like the macro pattern we've already described, but red **"Atypical Shape"** warnings have been applied to graphs that we believe FAIL TO SHOW the normal/characteristic pattern. To be "typical" (non-higher-ed) sites should exhibit (a) an initial stable low plateau, (b) a later persistent higher plateau, (c) an abrupt transition from the lower level to the newer level, and (d) once the new level was attained, it should not begin to recede ("the new level should be flat," not "hill shaped"). For higher ed sites, the "typical" shape is one that gradually rises to a "hill" then descends to a new value somewhere between the original level and the passed peak (e.g., the initial region may gradually ramp, and the new "plateau" tends not to be maintained).

Reported TTLs are from the live DNS for the final FQDN.

When multiple TTLs are mentioned, two different name servers reported two different TTL values.

The reported TTLs are from the time this report was prepared, and may have been different previously or subsequently. We include them here to demonstrate that the observed volumetric pattern does not appear to be affected by the TTL for a particular RRname/RRtype combination. (For example 3a) and 3b) both exhibit the "same shape" even though one has a TTL of 172800 seconds (two days) while the other has a TTL of just 60 seconds.)

Figure 3. Daily Query Count For Selected FQDNs From *.forbes.com
Split Out By Name and RRType

a) forbes.com NS (TTL=172800)

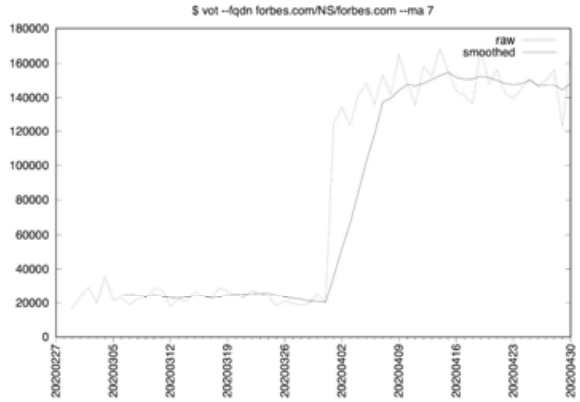

b) fast.forbes.com CNAME (TTL=60)

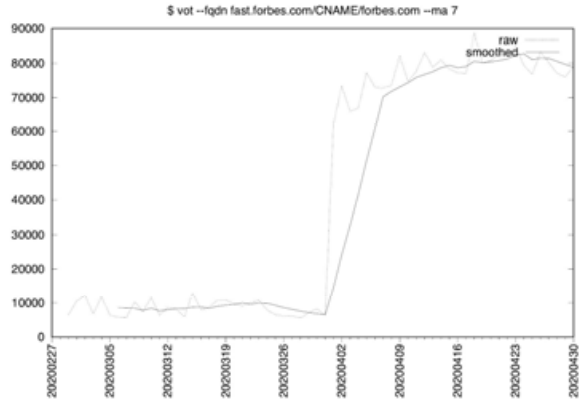

c) fuse.forbes.com CNAME (TTL=300)

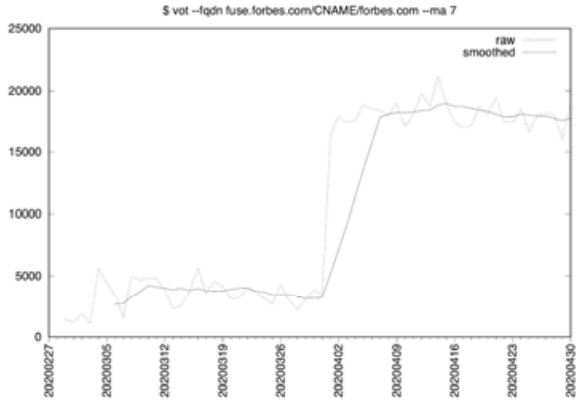

d) email.forbes.com NS (TTLs=300 and 172800)

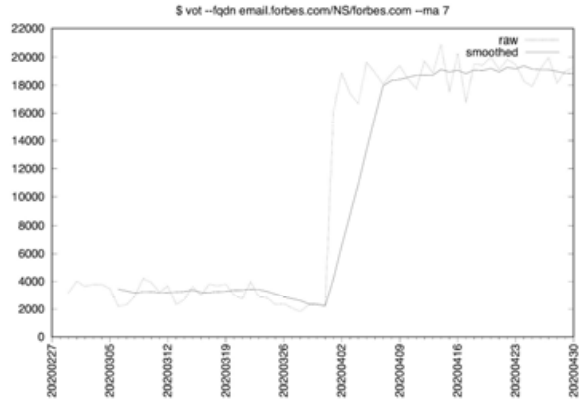

e) blogs-images.forbes.com CNAME (TTL=300)

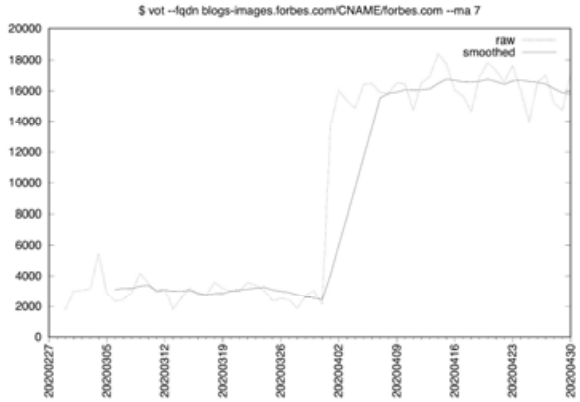

f) forbes.com SOA (TTL=900) [ATYPICAL SHAPE]

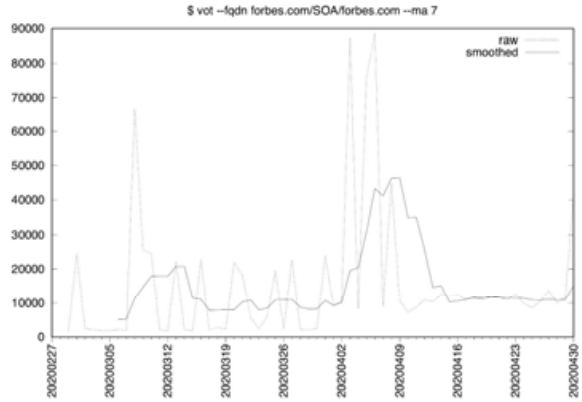

g) forbes.com A (TTL=300) [ATYPICAL SHAPE]

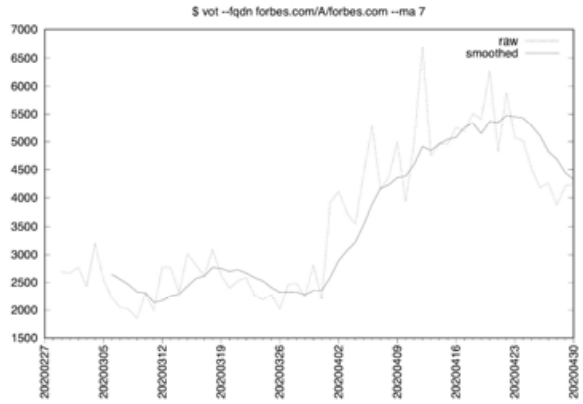

h) geolocation.forbes.com CNAME (TTL=3600)

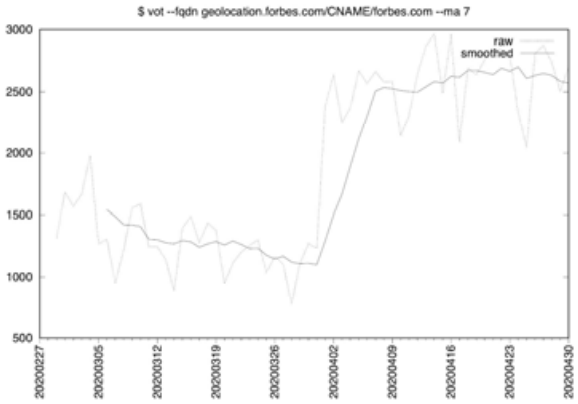

i) www.forbes.com CNAME (TTL=86400) [ATYPICAL SHAPE]

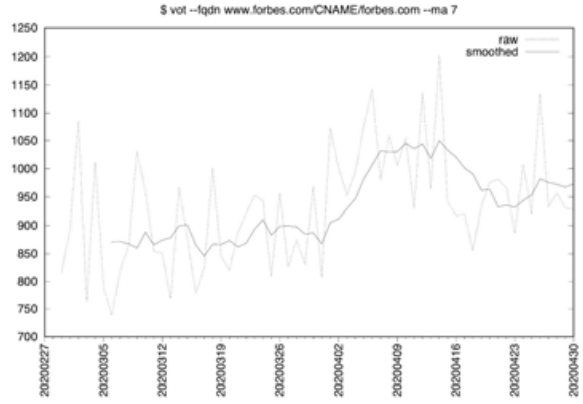

j) aax.forbes.com CNAME (TTL=300)

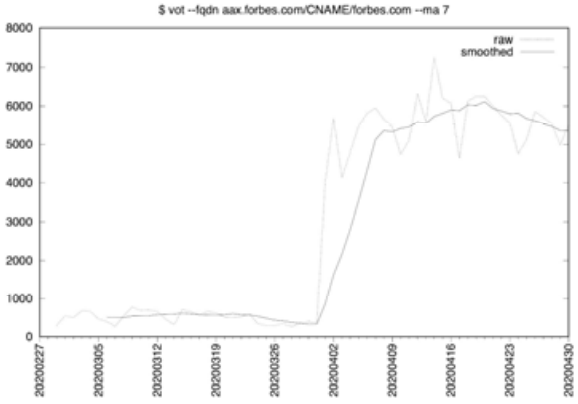

k) thumbor.forbes.com CNAME (TTL=86400) <-- [ATYPICAL SHAPES] --> l) images.forbes.com CNAME (TTL=300)

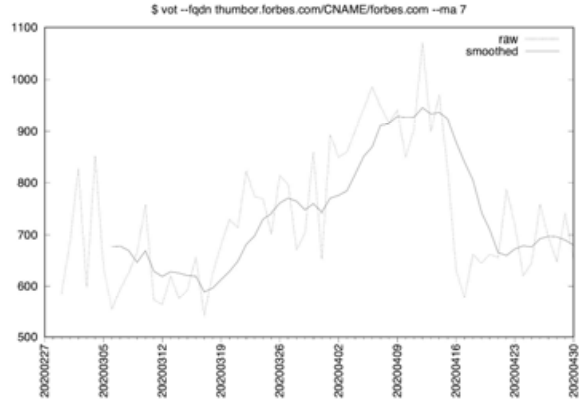

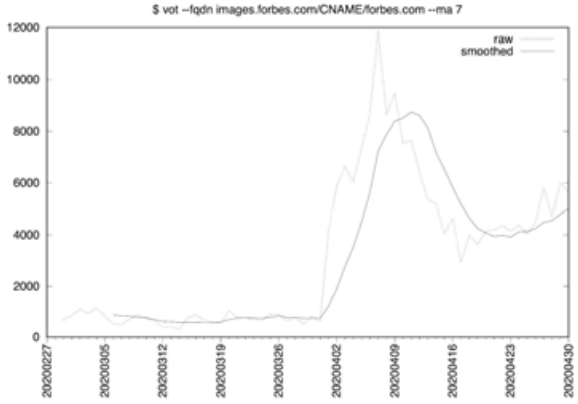

m) related.forbes.com CNAME (TTL=86400) <-- [ATYPICAL SHAPES] --> n) damapi.forbes.com NS (TTLs=300 and 21600)

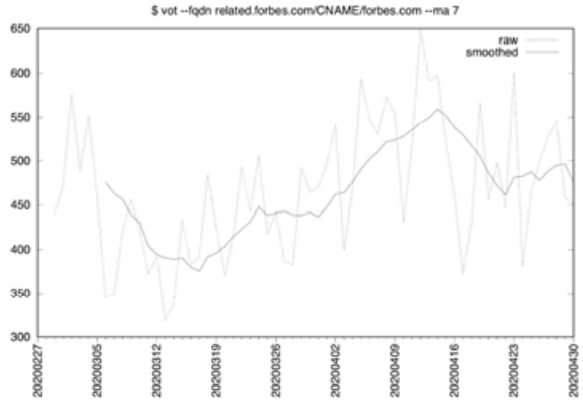

o) www3.forbes.com CNAME (TTL=21600)

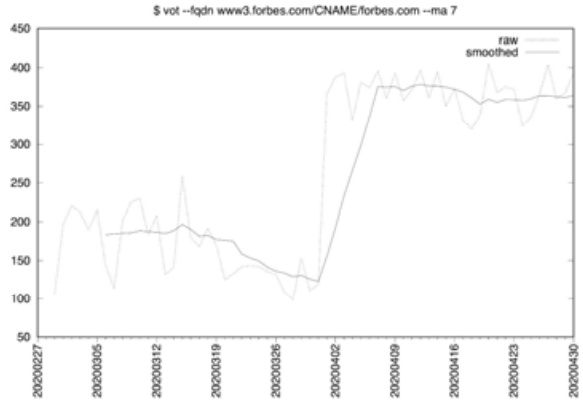

p) login.forbes.com A (TTL=300)

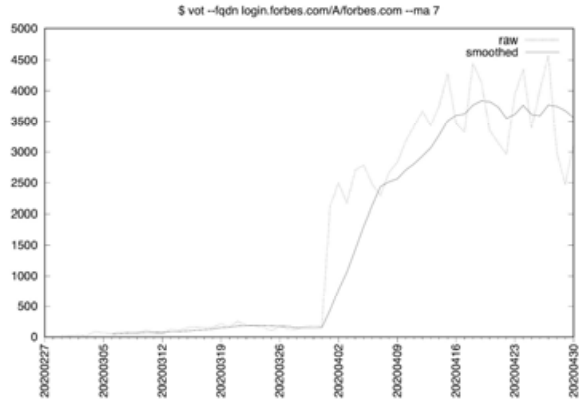

6) Anomalous Traffic

We saw "spikes" in volume for some sites, spikes that were so large that the spikes caused most of the "normal variation" in traffic volume to "wash out" due to the dominance of the spike or spikes. We think those spikes represent denial of service (DDoS) attack traffic reflexively targeting some unrelated third-party site.[12]

Let's consider three fairly blatant examples we observed: American Airlines, Netflix and Apple.

Example 1 -- American Airline's domain, aa.com, has an obvious spike near the end of our study period:

Figure 4. *.aa.com/any/aa.com query volume over time

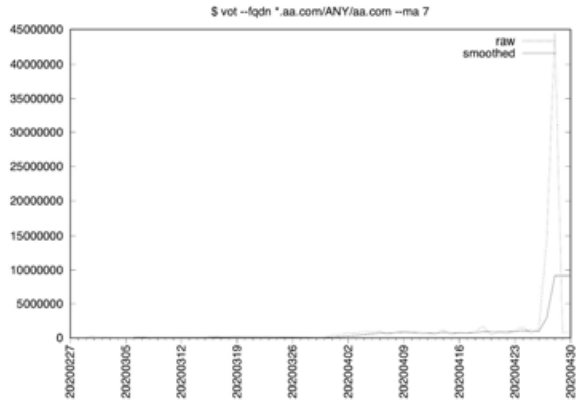

If we dig into the archived raw daily MTBL file for that date we can identify the RRset associated with that spike. The "spike day" looks like it occurred on **20200428**, so let's select that file for examination:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200428.D.mtbl
```

We'll then ask to see the highest count records for *.aa.com for that day:

```
$ dnstable_lookup -j rrset \*.aa.com. any aa.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | sort -nr | less
37317302 aa.com. SOA ["asia3.akam.net. hostmaster.akamai.com. 2013042368 9600 3600 604800
900"]
[...]
```

That's over **37.3 million Start of Authority (SOA)[13] queries for aa.com**!

That's just a **TON** of SOA record traffic! "Pretty printed," one specific RRset looks like:

```
$ dnstable_lookup -j rrset aa.com. soa aa.com | jq -r '.' | less
[...]
{
  "count": 37317302,
  "time_first": 1588048495,        <-- Tue Apr 28 04:34:55 UTC 2020
  "time_last": 1588106782,         <-- Tue Apr 28 20:46:22 UTC 2020
  "rrname": "aa.com.",
```

[12] For an approachable introduction to various DDoS-related vocabulary items, see "Spotting a Denial of Service (DoS) Attack," https://www.farsightsecurity.com/assets/media/infographics/ddos-infographic.pdf

[13] https://en.wikipedia.org/wiki/SOA_record

```
  "rrtype": "SOA",
  "bailiwick": "aa.com.",
  "rdata": [
    "asia3.akam.net. hostmaster.akamai.com. 2013042368 9600 3600 604800 900"
  ]
}
```

For comparative purposes, if we check some other date, such **20200320**, we see a much more typical value for the domain's SOA:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200320.D.mtbl
$ dnstable_lookup -j rrset aa.com. any aa.com | jq '.' | less
[...]
{
  "count": 54564,                 <-- the Apr 28th count is 683x this more typical value!
  "time_first": 1584644885,       <-- Thu Mar 19 19:08:05 UTC 2020
  "time_last": 1584716160,        <-- Fri Mar 20 14:56:00 UTC 2020
  "rrname": "aa.com.",
  "rrtype": "SOA",
  "bailiwick": "aa.com.",
  "rdata": [
    "asia3.akam.net. hostmaster.akamai.com. 2013042296 9600 3600 604800 900"
  ]
}
```

Example 2 -- Netflix.com: Another Example of A Dramatic SOA-driven Spike in Traffic: Another example of a major spike associated with extraordinary query volumes can be seen in the case of Figure 6 (notice the spike 3/4ths of the way across):

Figure 6. *.netflix.com/any/netflix.com query volume over time

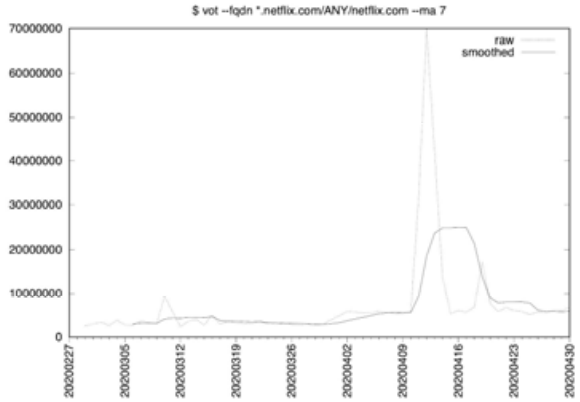

Checking that, we see:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200419.D.mtbl
$ dnstable_lookup -j rrset \*.netflix.com any netflix.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | sort -nr > netflix.txt
$ wc -l netflix.txt
906216 netf.txt           <-- that's a LOT of unique RRsets for just one day
```

Let's check to see what the distribution of record types looks like:

```
$ dnstable_lookup -j rrset \*.netflix.com any netflix.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | awk '{print $3}' | sort | \
```

```
uniq -c | sort -nr
 493482 TXT
 412307 CNAME
    260 A
    136 AAAA
     26 NS
      4 MX
      1 SOA
```

If you were casually looking at that distribution, you might think that the bulk of the traffic we'd seen consisted of TXT and CNAME records, but that's not the case -- the raw numbers of records aren't weighted by their counts. Let's total up the counts for each record type. For example, for SOA's:

```
$ dnstable_lookup -j rrset \*.netflix.com any netflix.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | grep " SOA " | \
awk '{print $1}' | paste -sd+ - | bc
10249514
```

Repeating that process for the other record types, we see a completely different picture overall:

Record Type	Count
SOA	10,249,514
NS	3,959,052
CNAME	3,102,747
TXT	997,714
A	398,769
AAAA	169,253
MX	3,395

```
$ more netflix.txt
10249514 netflix.com. SOA ["ns-81.awsdns-10.com. awsdns-hostmaster.amazon.com. 1
 7200 900 1209600 1800"]  <-- that's a ton of queries for the domain's SOA record
3681526 netflix.com. NS ["ns-81.awsdns-10.com.","ns-659.awsdns-18.net.","ns-1372
.awsdns-43.org.","ns-1984.awsdns-56.co.uk."]
[...]
```

This example also appears to demonstrate abuse of a wildcard pointing at obiwan-wc.geo.netflix.com given the following:

```
$ dnstable_lookup -j rrset \*.netflix.com any netflix.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | grep " CNAME " | \
grep "obiwan-wc.geo.netflix.com" | sort -nr > netflix-recs.txt
$ wc -l netflix-recs.txt
380032
$ awk '{print $1}' < netflix-recs.txt | paste -sd+ - | bc
687977
$ less netflix-recs.txt
1162 opstools.obiwan.netflix.com. CNAME ["obiwan-wc.geo.netflix.com."]
19 uswest.obiwan.netflix.com. CNAME ["obiwan-wc.geo.netflix.com."]
11 www.www.www.www.www.www.www.www.www.www.www.www.www.assets.obiwan.netflix.com. CNAME
["obiwan-wc.geo.netflix.com."]
11 www.www.www.www.www.www.www.www.www.cms.obiwan.netflix.com. CNAME ["obiwan-
wc.geo.netflix.com."]
11 www.www.www.www.www.www.www.uswest.obiwan.netflix.com. CNAME ["obiwan-
wc.geo.netflix.com."]
11 www.www.www.www.cms.obiwan.netflix.com. CNAME ["obiwan-wc.geo.netflix.com."]
[etc]
```

Example 3 -- Apple: Another interesting example of anomalous traffic can be seen the graph of *.apple.com volume (notice the spike near the right-hand-side of the graph):

Figure 7. *.apple.com/any/apple.com query volume over time

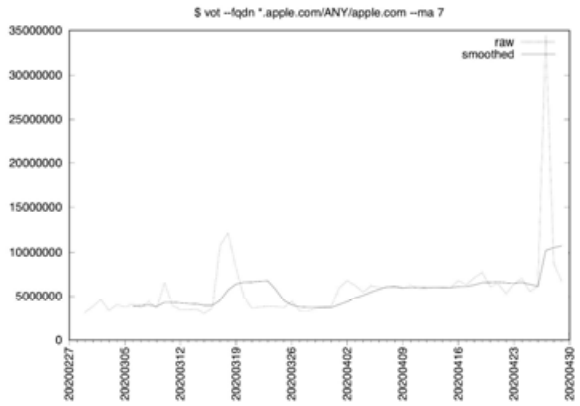

If we disect the spike, we see over 18.5 million SOA queries in one day's results, obviously a substantial "contribution" to that spike:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200428.D.mtbl
$ dnstable_lookup -j rrset \*.apple.com. soa apple.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | sort -nr | \
awk '{print $1}' | paste -sd+ - | bc
18532290
```

For comparison, a more-normal day's SOA volume for *.apple.com looks like:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200315.D.mtbl
$ dnstable_lookup -j rrset \*.apple.com. any apple.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | grep "SOA" | \
awk '{print $1}' | paste -sd+ - | bc
766933    <-- the 20200428 data is over 24X this level!
```

Checking in more detail, the top SOA records associated with that spike from the 20200428 data looked like:

```
$ export DNSTABLE_FNAME=/export/dnsdb/mtbl/dns.20200428.D.mtbl
$ dnstable_lookup -j rrset \*.apple.com. soa apple.com | jq -r '"\(.count) \(.rrname)
\(.rrtype) \(.rdata)"' | sort -nr | less
17311824 apple.com. SOA ["nserver.apple.com. hostmaster.apple.com. 2010115035 900 900
2016000 14400"]
411557 euro.apple.com. SOA ["nserver.apple.com. hostmaster.euro.apple.com. 2010092407 900
900 604800 86400"]
341904 asia.apple.com. SOA ["nserver.apple.com. hostmaster.apple.com. 2012021895 900 900
2592000 1800"]
251296 health.apple.com. SOA ["nserver.apple.com. hostmaster.apple.com. 20 1800 900
2592000 1800"]
50926 apple.com. SOA ["nserver.apple.com. hostmaster.apple.com. 2010115037 900 900
2016000 14400"]
34564 apple.com. SOA ["nserver.apple.com. hostmaster.apple.com. 2010115036 900 900
2016000 14400"]
[all remaining SOA's have counts <30,000]
```

There MUST have been additional record types that were ALSO part of that spike. Let's check the "A" records:"

```
$ dnstable_lookup -j rrset \*.apple.com. any apple.com | jq -r '"\(.count) \(.rrname)
\(.rrtype) \(.rdata)"' | grep " A " | awk '{print $1}' | paste -sd+ - | bc
6016605       <-- For comparison, the 20200315 data was just 303,608
```

We'll repeat that process for other common record types. Summarizing what we saw for a variety of selected RRtype values:

Table II. Relative volume for *.apple.com for selected record types for two dates

Record Type	Count (2020-04-28)	Count (2020-03-15)	Ratio
SOA	18,532,290	766,933	24.16
A	6,016,605	303,608	19.81
CNAME	4,364,149	2,219,937	1.96
AAAA	4,040,620	173,706	23.26
TXT	2,459,745	5,620	**437.67**
NS	1,508,618	81,471	18.51

The volume of TXT records particularly attracted our attention since it was over **437x** normal levels (dang!). Eyeballing the TXT records, we noticed that **95% of those TXT records were SPF-related:**

```
$ dnstable_lookup -j rrset \*.apple.com. any apple.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | grep " TXT " | \
grep "_spf" | awk '{print $1}' | paste -sd+ - | bc
2350855            <-- 2350855/2459745*100=95.00%
```

```
$ dnstable_lookup -j rrset \*.apple.com. any apple.com | \
jq -r '"\(.count) \(.rrname) \(.rrtype) \(.rdata)"' | grep " TXT " | \
grep "_spf" | less
49037 push.apple.com. TXT ["\"count=50\"","\"v=spf1 include:_spf.apple.com include:_spf-
txn.apple.com ~all\""]
1457 _spf-txn.apple.com. TXT ["\"v=spf1 ip4:17.151.1.0/24 ip4:17.171.37.0/24
ip4:17.111.110.0/23 ~all\""]
1012 email.apple.com. TXT ["\"v=spf1 include:_spf-txn.apple.com include:_spf-
mkt.apple.com include:_spf.apple.com ~all\""]
831 _spf-mkt.apple.com. TXT ["\"v=spf1 ip4:17.171.23.0/24 ip4:17.179.250.0/24
ip4:17.32.227.0/24 ~all\""]
741 _spf.apple.com. TXT ["\"v=spf1 ip4:17.151.62.66 ip4:17.151.62.67 ip4:17.151.62.68
ip4:17.171.2.60 ip4:17.171.2.68 ip4:17.171.2.72 ~all\""]
[...]
1 account.apple.com. TXT ["\"v=spf1 redirect=_spf.apple.com\""]
```

Checking with dig and a random hostname that we made up ourselves, we can see that apple.com's name servers are apparently set up to do **wildcarding**, responding to **<anything>**.apple.com (at least for TXT record queries). Any such query routinely reports an ***SPF redirect*** to _spf.apple.com:

```
$ dig <randomtexthere>.apple.com txt +short
"v=spf1 redirect=_spf.apple.com"
```

We believe this may be getting exploited for pseudo-random subdomain DDOS attack purposes.

Individual sites and vendors of authoritative name servers should ensure those servers are configured to do Response Rate Limiting (RRL).[14] Unfortunately, at least as of a 2017 talk that Casey Deccio did, **only about 17% of Internet authoritative name servers use RRL.**[15]

[14] https://kb.isc.org/docs/aa-01000

[15] https://indico.dns-oarc.net/event/27/contributions/462/attachments/462/764/2017-09-29-rrl-oarc.pdf at slide 17.

Section II. Graphs for the 316 Sites

Because most readers won't have access to the archived daily data files the way we do, we wanted to share a selection of volume over time graphs for a variety of different sites that we expected might be impacted by the coronavirus pandemic.

News and Partisan Sites

During the pandemic, individuals forced to remain at home have been eager[16] for news about topics such as:

- The coronavirus itself, including testing information and disease statistics
- Work on potential vaccines and treatments for the disease
- Stay-at-home orders and news about when various businesses or areas may reopen
- Economic impacts and support programs (e.g., the Paycheck Protection Program, food banks, etc.)

That said, how you view the pandemic and its impacts may be shaped by your political orientation, with substantial differences in perception between those on the left and those on the right.[17]

We therefore wanted to see if traffic associated with Internet news sites was up across the political spectrum, or moreso for partisan liberal/left leaning news sites, for conservative/right leaning news sites, or for both. (Differences in this respect was a major theme in the New York Times story).

A variety of sources provide opinions on media bias today.[18] Some of those rating sites cases decouple bias in "hard news" reporting from bias in opinion pieces (e.g., for example, a site might offer relatively balanced "hard news" but consistently present highly partisan editorials favoring either the left or the right). We've elected to assess the bias of each site *as a whole*, in part because in online formats there normally aren't separate news and editorial "sections" -- content tends to all run together.

We have not formally adopted any particular site's exact bias categorization, but needless to say, if you disagree with our assignments, you should feel free to reassign sites to your preferred category (it won't fundamentally affect the results reported here -- virtually all news sites have exhibited an increase in DNS traffic volume).

We also note for the record that the inclusion (or exclusion) of any site in this section (or any other section) is not meant as an endorsement (or indication of disapproval) of any site by Farsight.

Some sites, such as some news aggregation sites, are included because of their acknowledged disproportionate influence.[19]

Other sites, what some might term "editorial" or "opinion sites" (such as Sean Hannity's), were included because they similarly had outsize influence -- reportedly not just on the public, but even on the President himself.[20]

[16] "The Virus Changed the Way We Internet," https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html

[17] "Americans divided on party lines over risk from coronavirus: Reuters/Ipsos poll," https://www.reuters.com/article/us-health-coronavirus-usa-polarization/americans-divided-on-party-lines-over-risk-from-coronavirus-reuters-ipsos-poll-idUSKBN20T2O3

[18] See for example https://www.allsides.com/media-bias/media-bias-ratings and https://www.adfontesmedia.com/

[19] "Drudge Report: Small Operation, Large Influence," https://www.journalism.org/2011/05/09/drudge-report-small-operation-large-influence/

[20] "Sean Hannity is the most influential TV host," https://www.washingtonpost.com/opinions/2019/07/26/sean-hannity-is-most-influential-tv-host/

1) News-Related Domains: Liberal/Left-Leaning

1. abcnews.com
2. bbc.co.uk
3. bloomberg.com
4. cbsnews.com
5. cnbc.com
6. cnn.com
7. dailymail.co.uk
8. guardian.co.uk
9. huffingtonpost.com
10. motherjones.com
11. msnbc.com
12. nbcnews.com
13. nytimes.com
14. pbs.org
15. politico.com
16. rollingstone.com
17. salon.com
18. sfgate.com
19. slate.com
20. telegraph.co.uk
21. theatlantic.com
22. thehill.com
23. usatoday.com
24. usnews.com
25. washingtonpost.com

1. ABC News

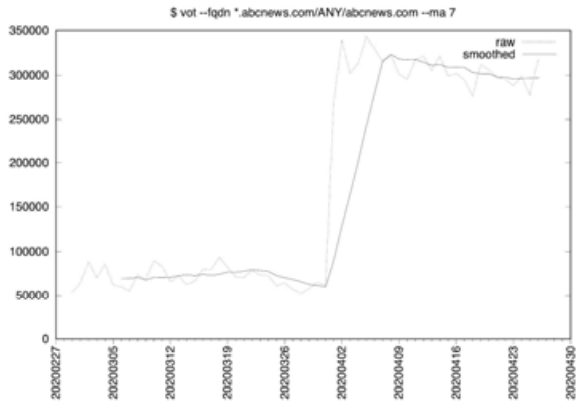

2. BBC [ATYPICAL SHAPE]

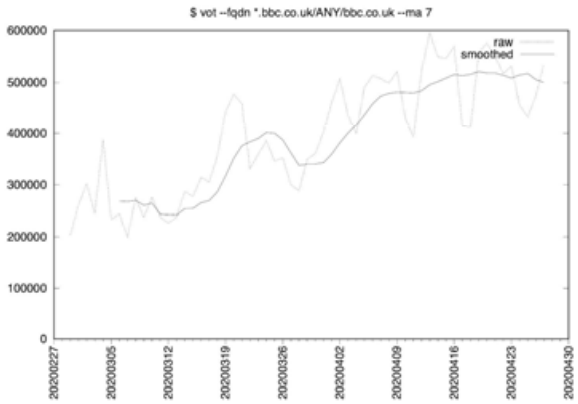

3. Bloomberg [ATYPICAL SHAPE]

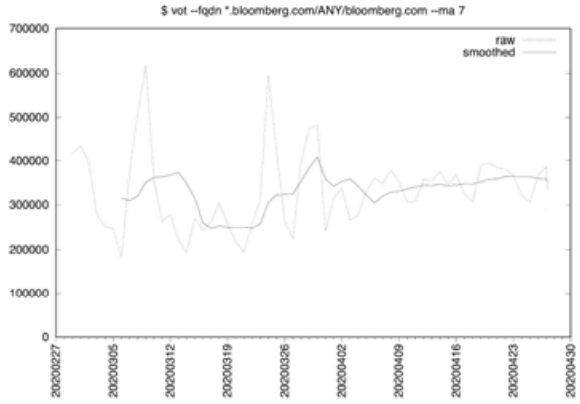

4. CBS News

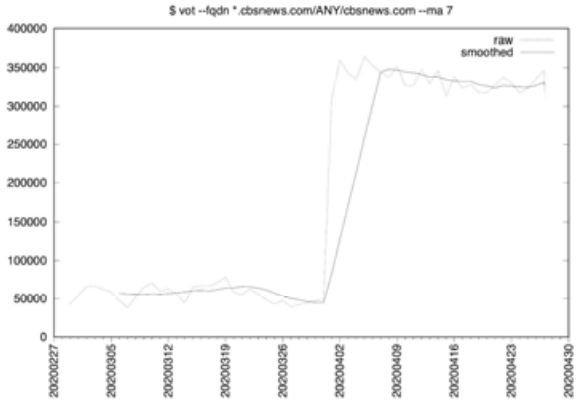

5. CNBC

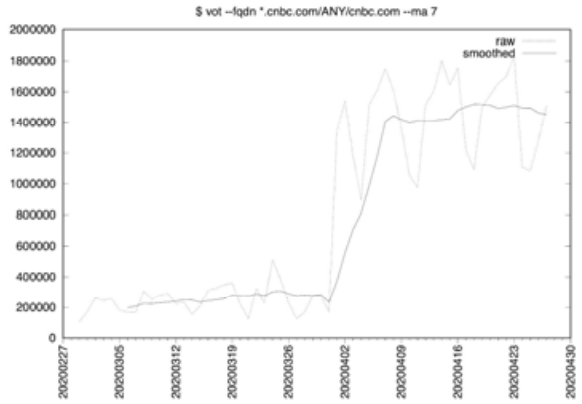

6. CNN

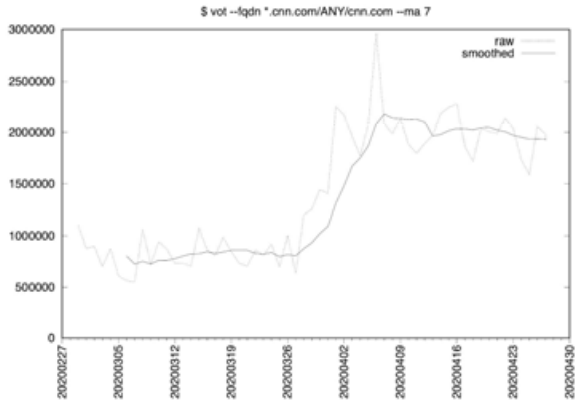

7. Daily Mail, England

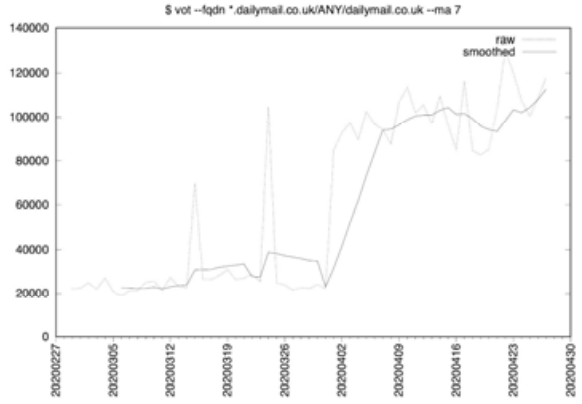

8. Guardian, England [ATYPICAL SHAPE]

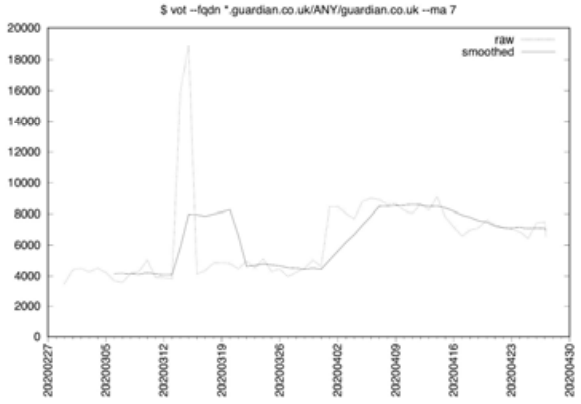

9. Huffington Post [ATYPICAL SHAPE]

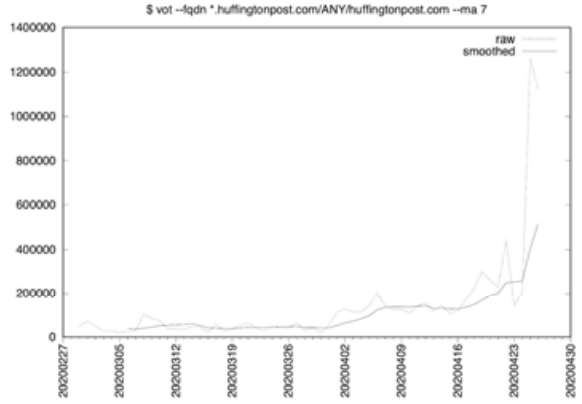

10. Mother Jones

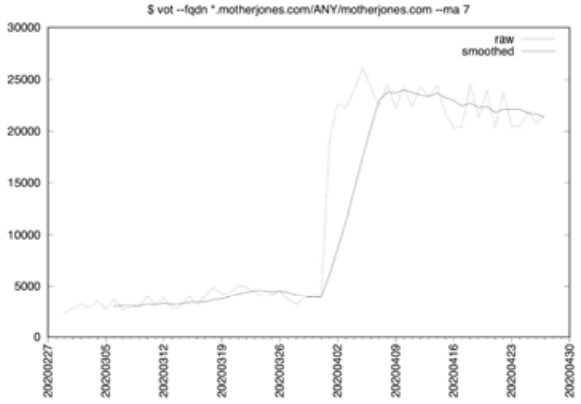

11. MSNBC

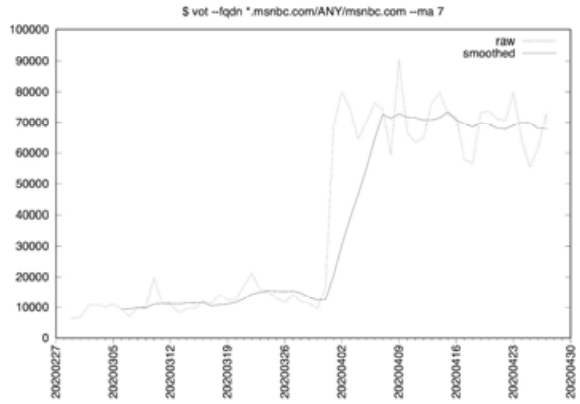

12. NBC News

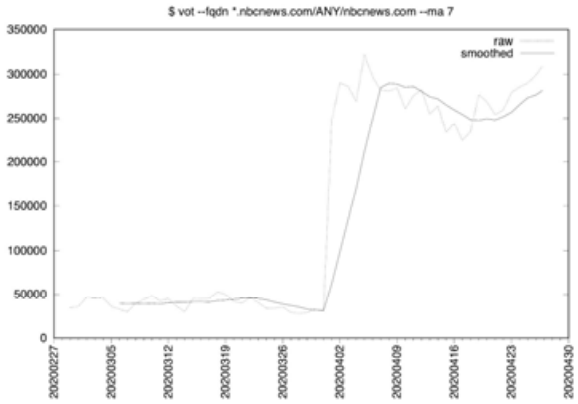

13. New York Times [ATYPICAL SHAPE]

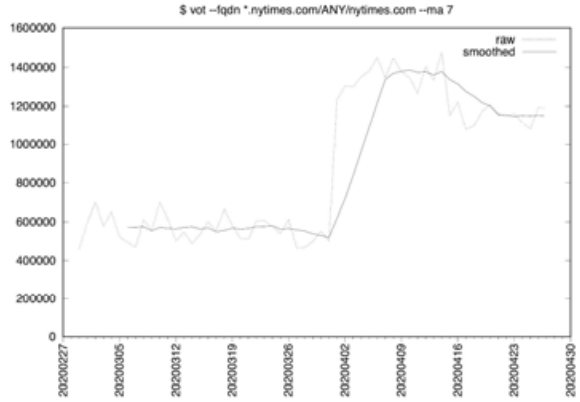

14. PBS [ATYPICAL SHAPE]

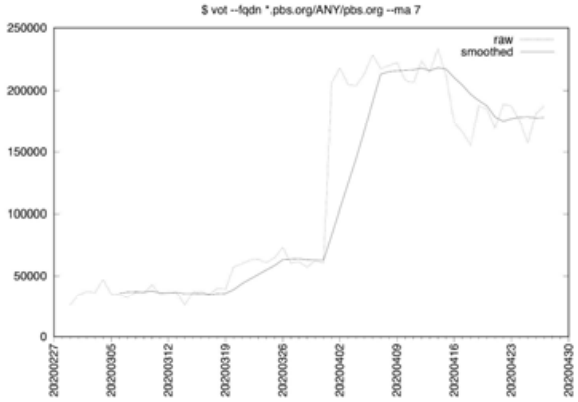

15. Politico [ATYPICAL SHAPE]

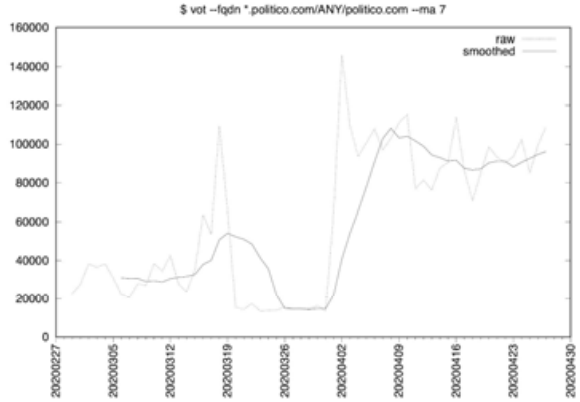

16. Rolling Stone

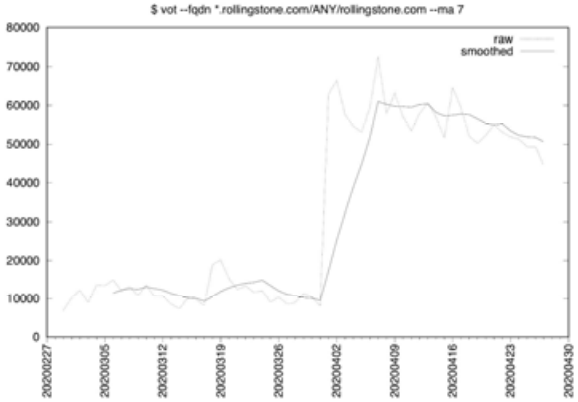

17. Salon

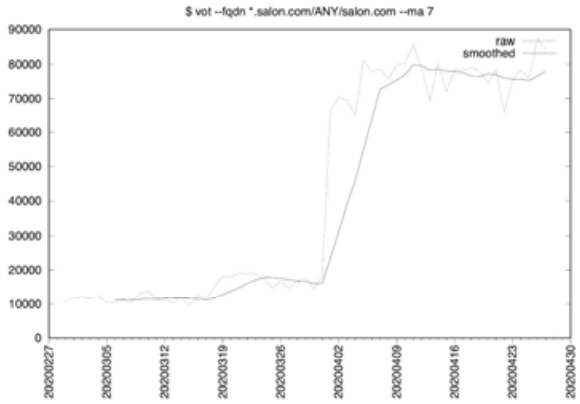

18. SFGate

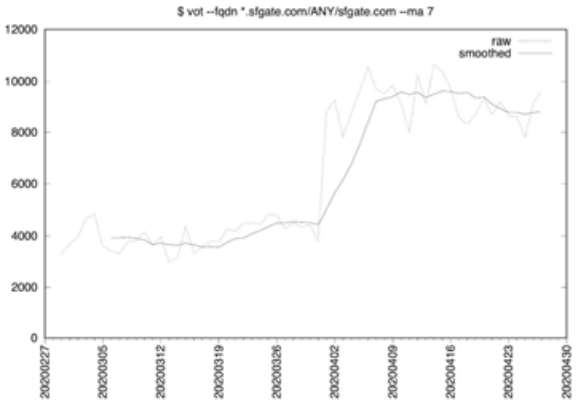

19. Slate

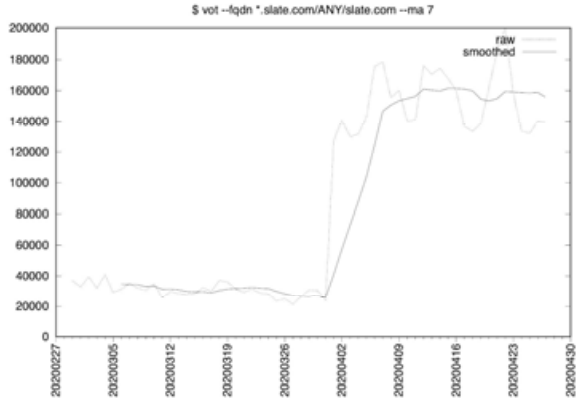

20. Telegraph, England

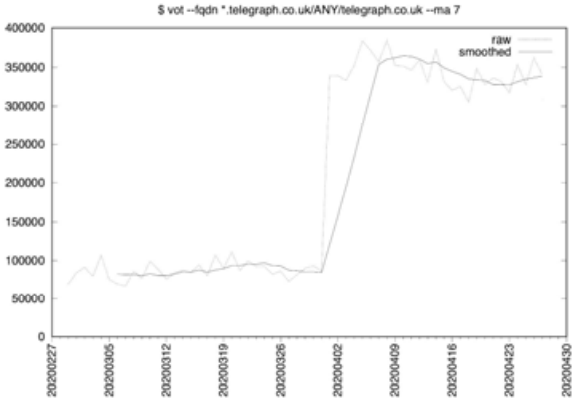

21. The Atlantic [ATYPICAL SHAPE]

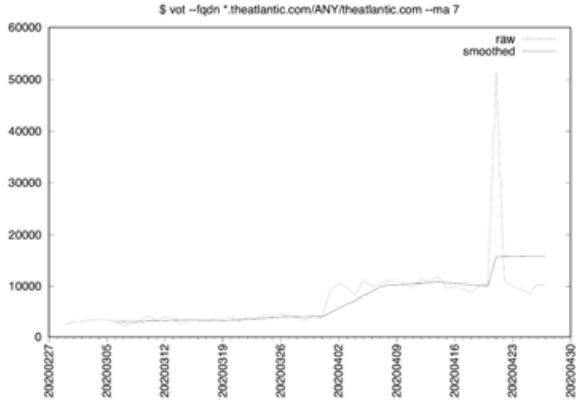

22. The Hill

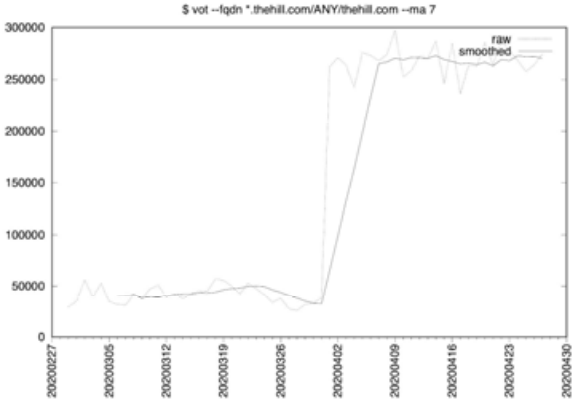

23. USA Today

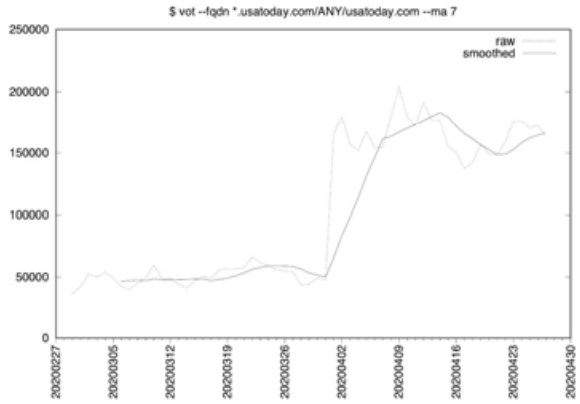

24. US News

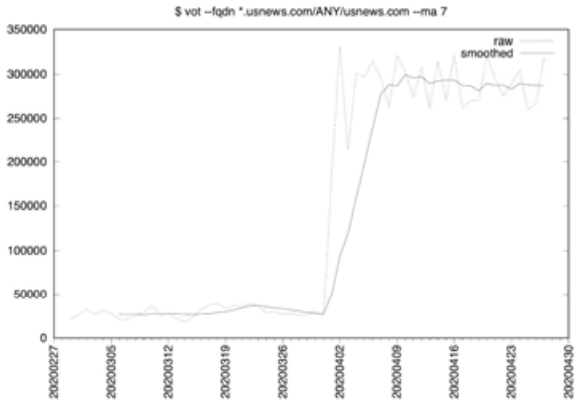

25. Washington Post

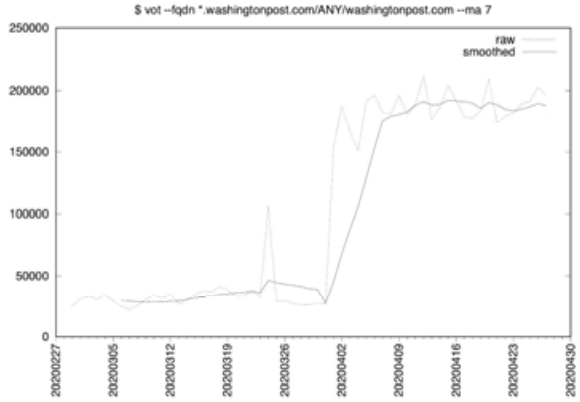

2) News-Related Domains: Neutral/Balanced

1. aljazeera.com
2. apnews.com
3. c-span.org
4. latimes.com
5. realclearpoliticscom
6. reuters.com
7. sky.com
8. upi.com

1. Aljazeera

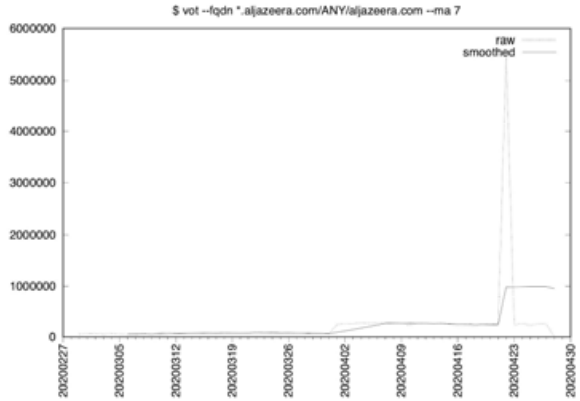

2. AP News

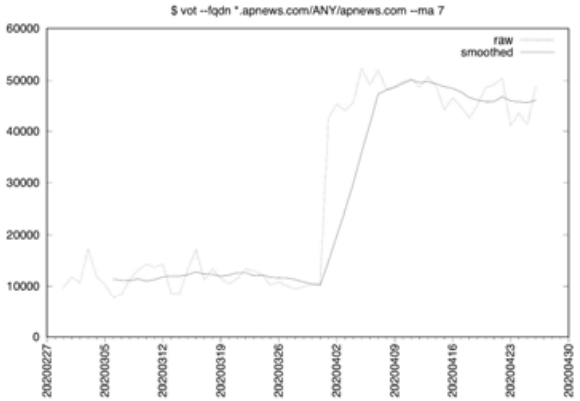

3. C-SPAN

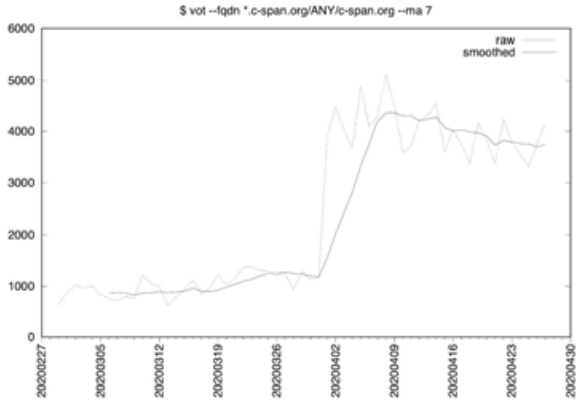

4. LA Times [ATYPICAL SHAPE]

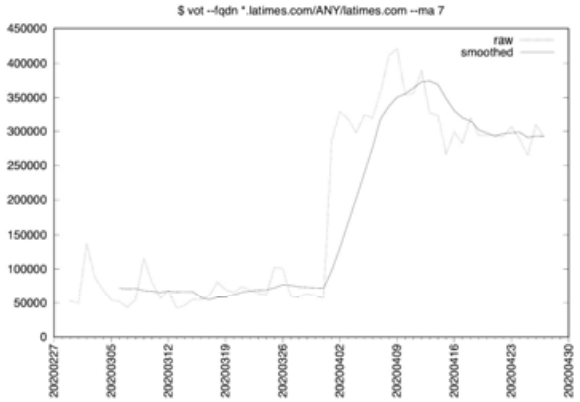

5. Real Clear Politics

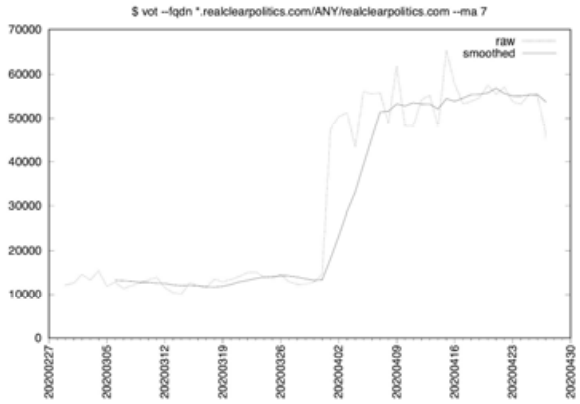

6. Reuters

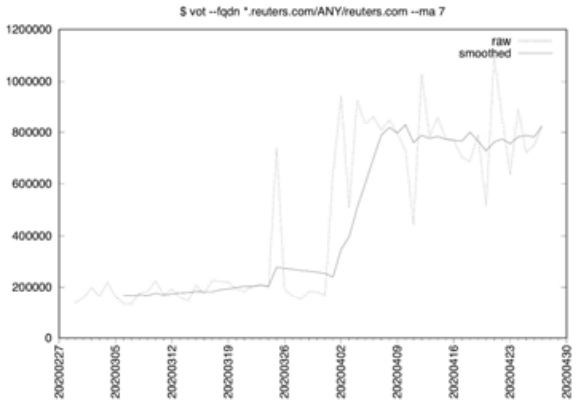

7. Sky News, England [ATYPICAL SHAPE]

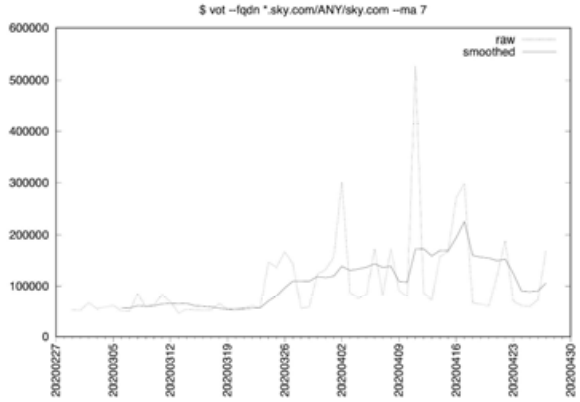

8. UPI [ATYPICAL SHAPE]

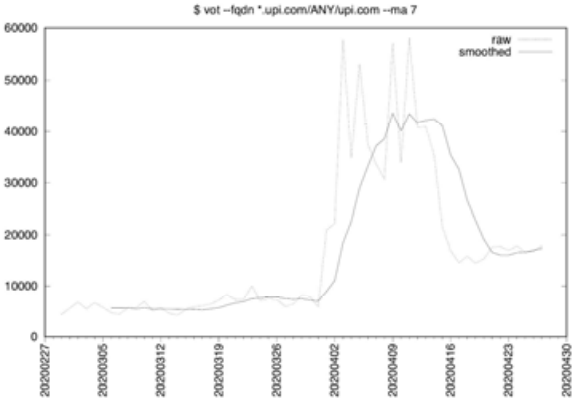

3) News-Related Domains: Conservative/Right-Leaning

1. breitbart.com
2. dailycaller.com
3. drudgereport.com
4. forbes.com
5. foxnews.com
6. hannity.com
7. lauraingraham.com
8. lauraloomer.us
9. michaelsavage.com
10. nationalreview.com
11. nypost.com
12. rushlimbaugh.com
13. thegatewaypundit.com
14. thesun.co.uk
15. washingtontimes.com
16. wsj.com
17. zerohedge.com

1. Breitbart

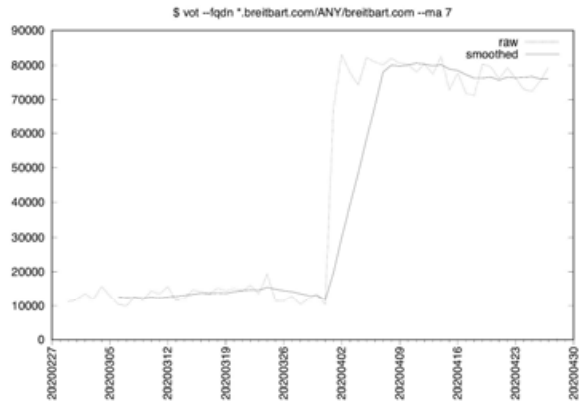

2. Daily Caller

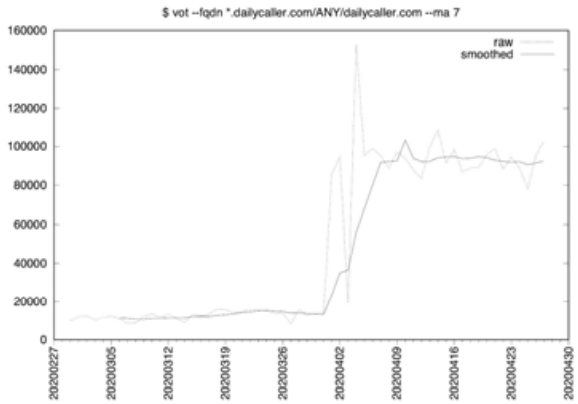

3. Drudge Report

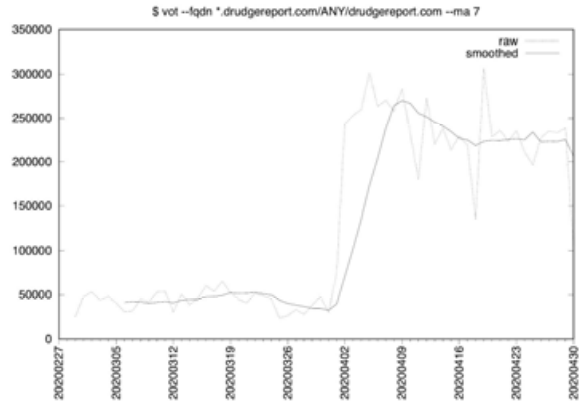

4. Forbes

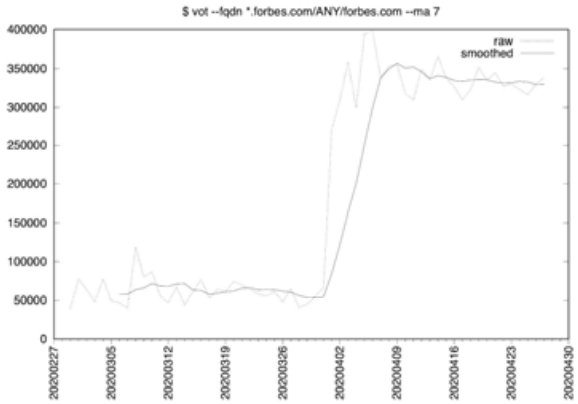

5. Fox News

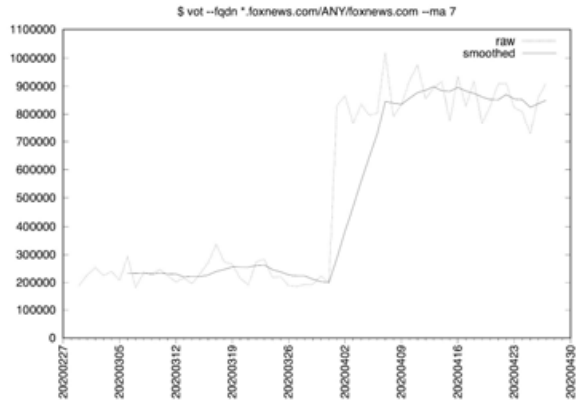

6. Hannity

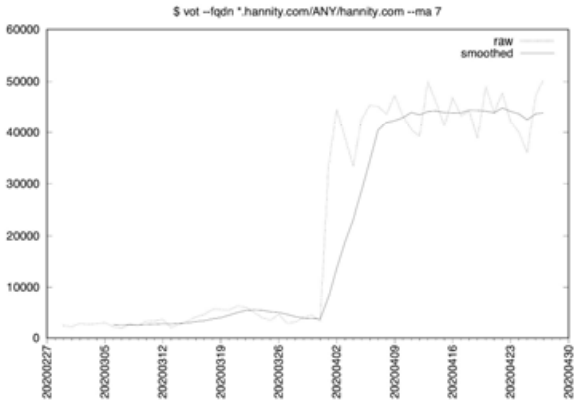

7. Laura Ingraham

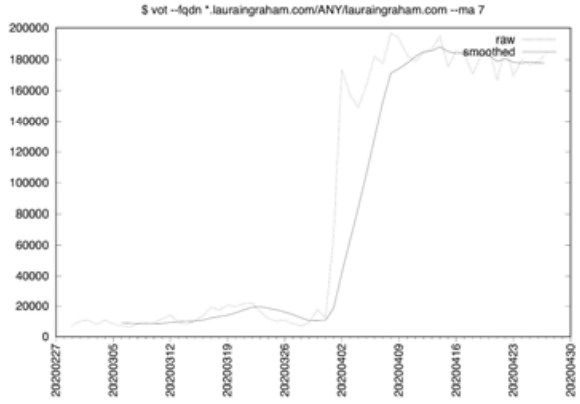

8. Laura Loomer

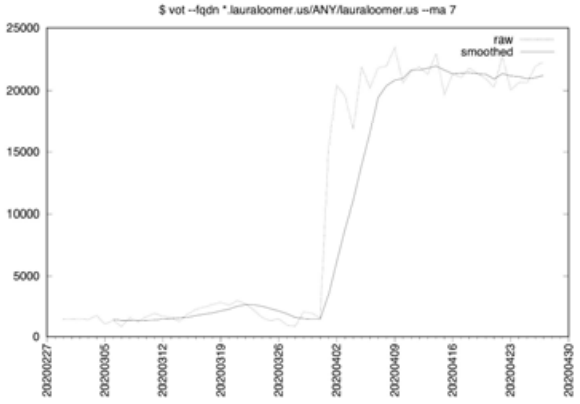

9. Michael Savage

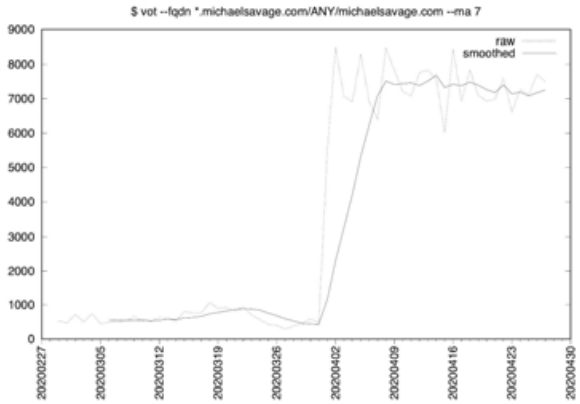

10. National Review

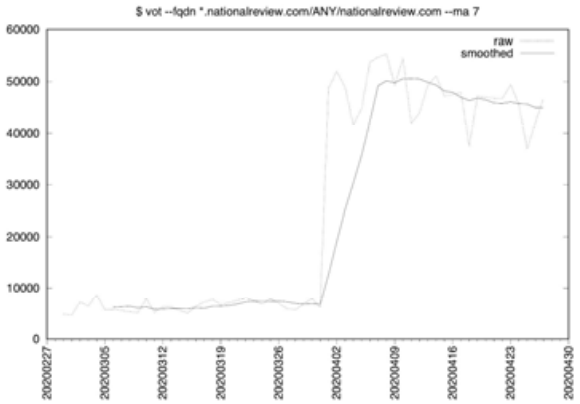

11. New York Post

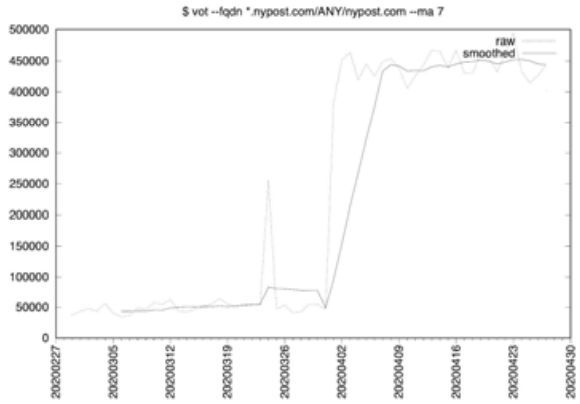

12. Rush Limbaugh [ATYPICAL SHAPE]

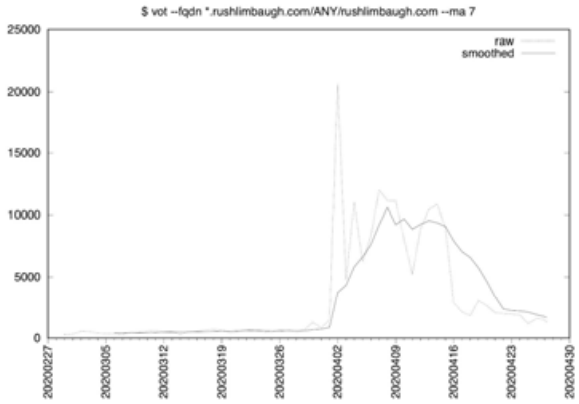

13. The Gateway Pundit (note "The" in this site's name)

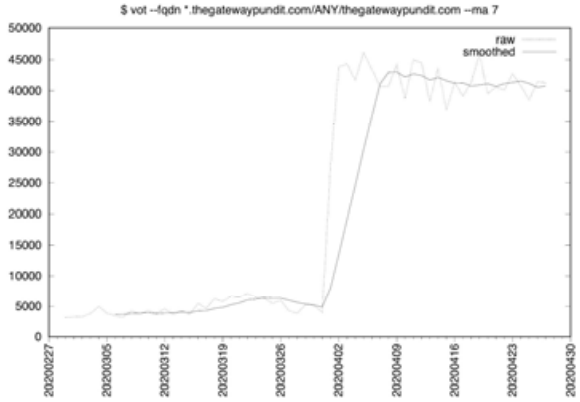

14. The Sun, England

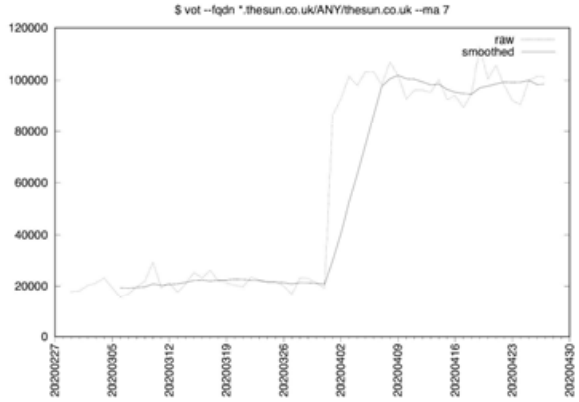

15. Washington Times

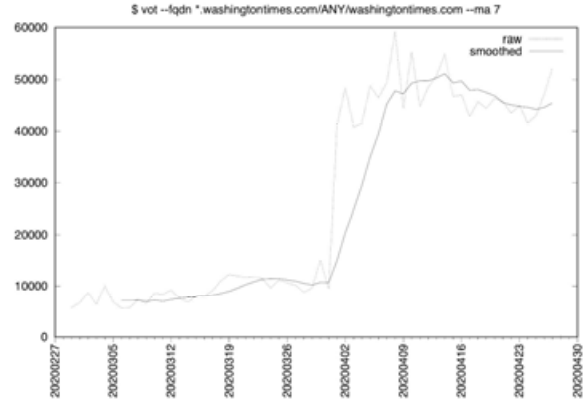

16. Wall Street Journal

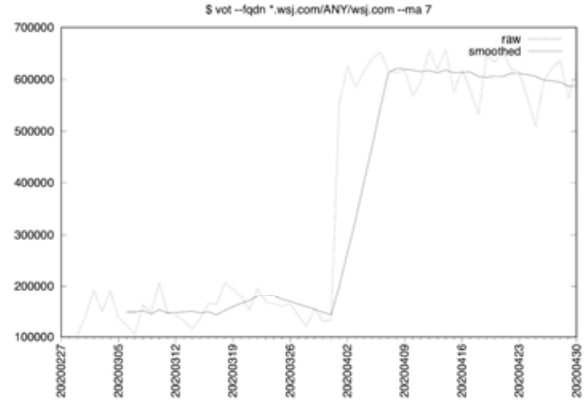

17. Zerohedge

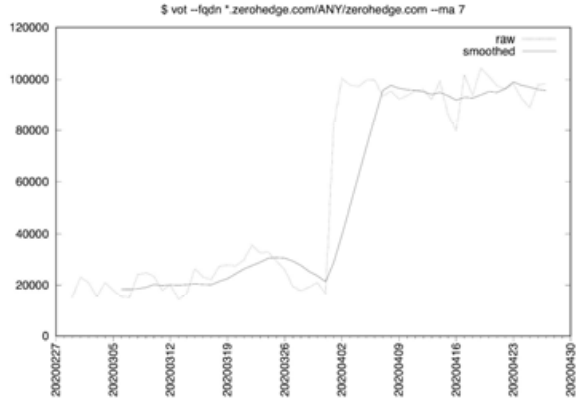

Travel and Transportation Sites

Widespread travel restrictions have affected business and leisure travel, including stay-at-home orders and border closures. Fewer people traveling means reduced demand for air travel, rental cars, hotels and restaurants, and demand for luxury cruise lines plummeted when cruise ship passengers were particularly hard hit by the virus.

At the same time, many people may have needed to cancel pre-scheduled travel and hoped to obtain refunds or credits, and demand for shipment of groceries and other staples remained strong -- people need to be able to feed their families.

We also recognized that some companies might have their staff members working remotely -- we might not be seeing JUST customer traffic in some cases.

Bottom line, we weren't sure in advance how DNS traffic levels for travel and transportation sites might be affected.

Having run the data, what we're seeing is more traffic in most cases, with some sites exhibiting spikes consistent with DDoS attacks exploiting those sites.

Airlines (alphabetized by domain)

1. aa.com
2. aerlingus.com
3. aeroflot.com
4. airasia.com
5. airchina.com
6. airfrance.com
7. airindia.com
8. airnewzealand.com
9. alaskaair.com
10. allegiantair.com
11. ana.co.jp
12. austrian.com
13. britishairways.com
14. brusselsairlines.com
15. cargolux.com
16. cathaypacific.com
17. ceair.com
18. cebupacificair.com
19. china-airlines.com
20. copaair.com
21. csair.com
22. delta.com
23. easyjet.com
24. egyptair.com
25. elal.com
26. emirates.com
27. ethiopianairlines.com
28. etihad.com
29. eurowings.com
30. evaair.com
31. flyasiana.com
32. flyfrontier.com
33. flylevel.com
34. flysas.com
35. flytap.com
36. garuda-indonesian.com
37. goair.in
38. goindigo.in
39. hainanairlines.com
40. hawaiianairlines.com
41. iberia.com
42. icelandair.com
43. iranair.com
44. jal.co.jp
45. jetblue.com
46. klm.com
47. koreanair.com
48. latam.com
49. lionair.co.id
50. lufthansa.com
51. malaysiaairlines.com
52. norwegian.com
53. philippineairlines.com
54. pobeda.aero
55. qantas.com
56. qatarairways.com
57. rossiya-airlines.com
58. ryanair.com
59. s7.ru
60. saudia.com
61. singaporeair.com
62. southwest.com
63. spicejet.com
64. spirit.com
65. swiss.com
66. thaiairways.com
67. transavia.com
68. turkishairlines.com
69. united.com
70. vietjetair.com
71. vietnamairlines.com
72. vueling.com
73. wizzair.com

American Airlines

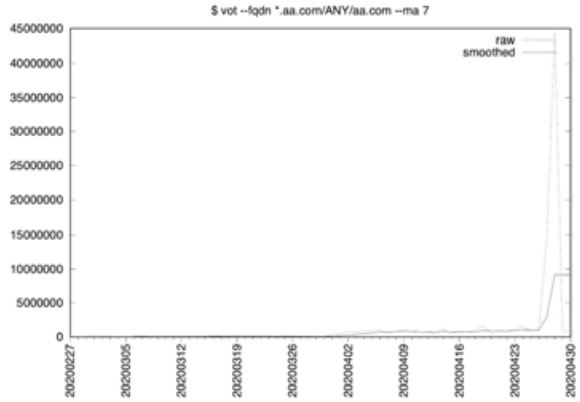

2. Aer Lingus

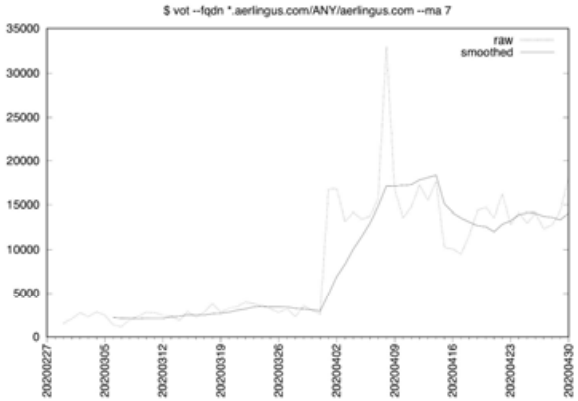

3. Aeroflot [ATYPICAL SHAPE]

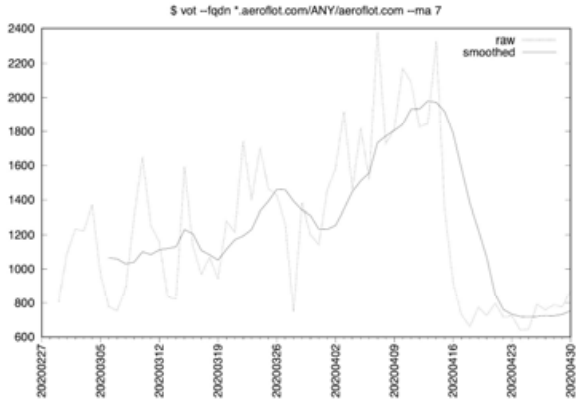

4. Air Asia [ATYPICAL SHAPE]

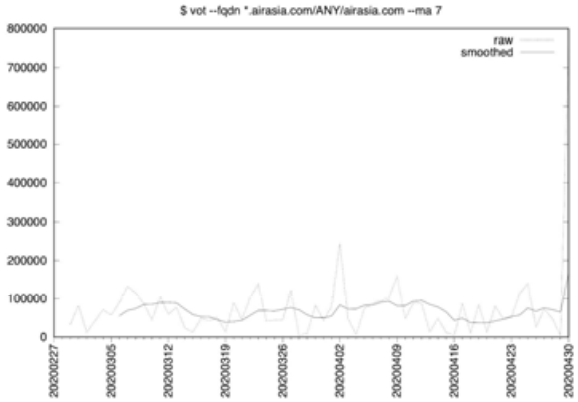

5. Air China [ATYPICAL SHAPE]

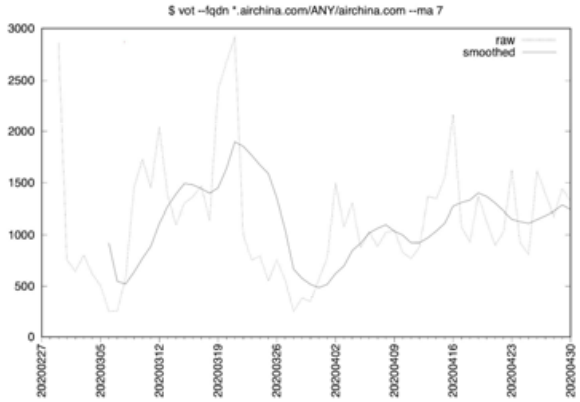

6. Air France

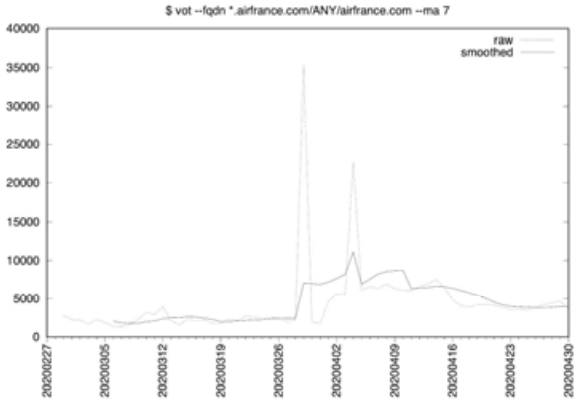

7. Air India [ATYPICAL SHAPE]

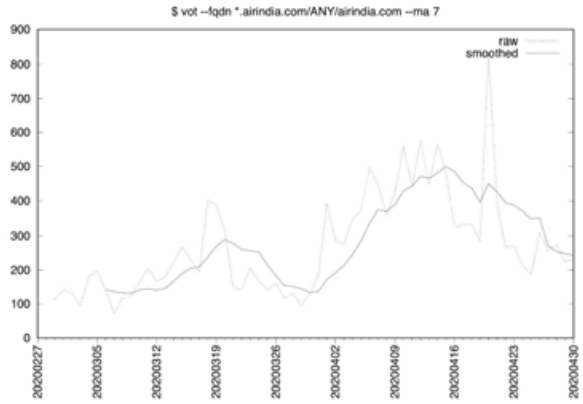

8. Air New Zealand

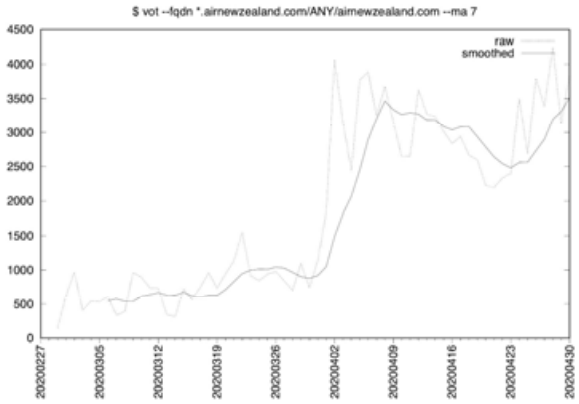

9. Alaska Airlines

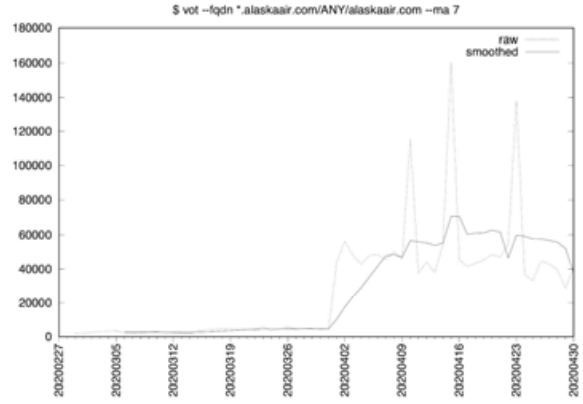

10. Allegiant Airlines

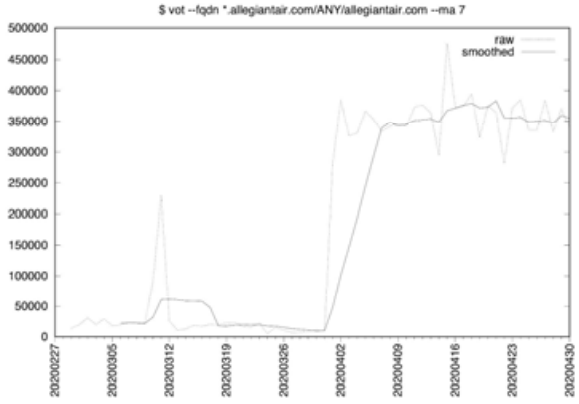

11. All Nippon Airways

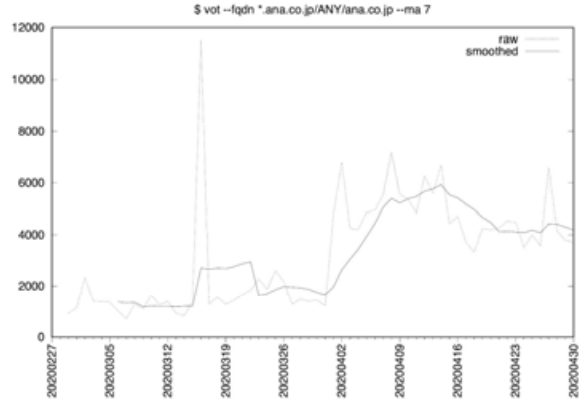

12. Austrian Airlines [ATYPICAL SHAPE]

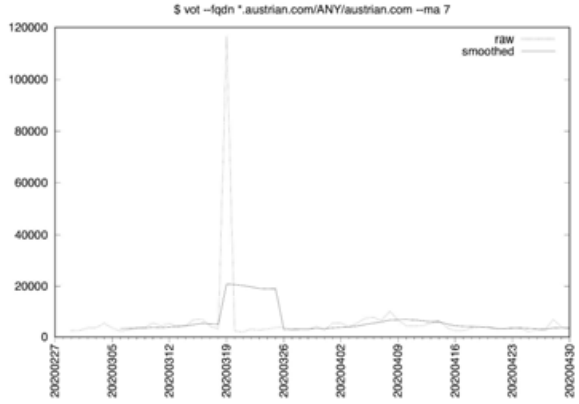

13. British Airways [ATYPICAL SHAPE]

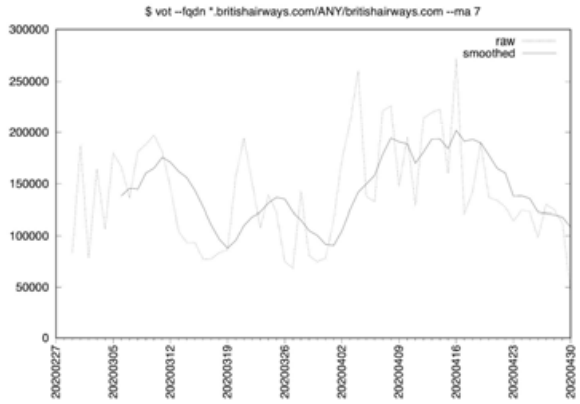

14. Brussels Airlines [ATYPICAL SHAPE]

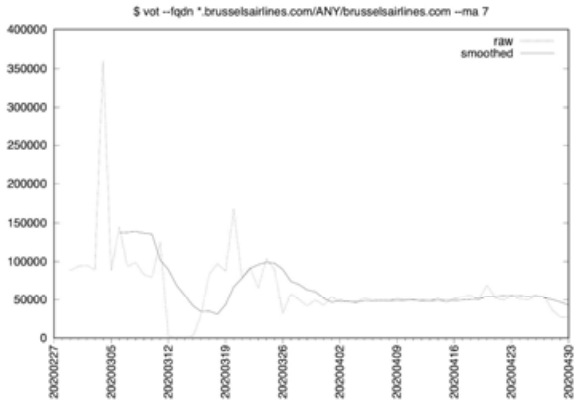

15. Cargolux Airlines [ATYPICAL SHAPE]

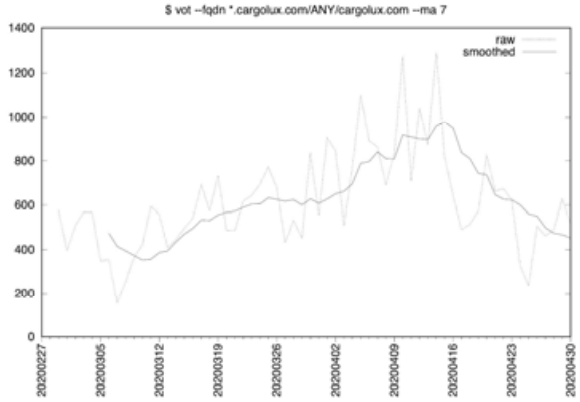

16. Cathay Pacific [ATYPICAL SHAPE]

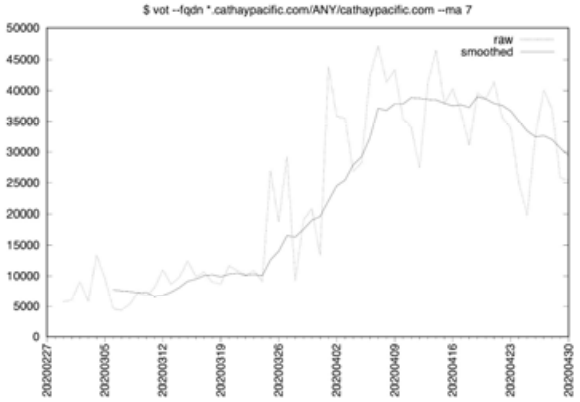

17. China Eastern Airlines [ATYPICAL SHAPE]

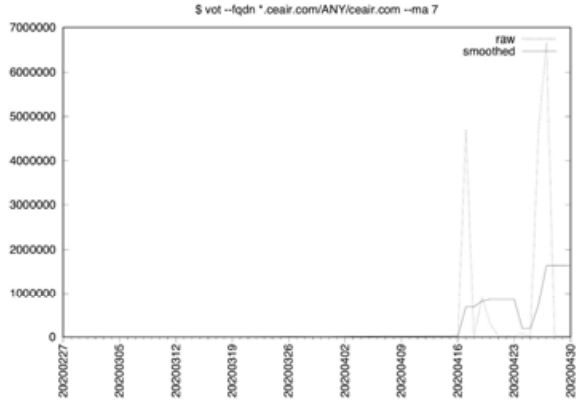

18. Cebu Pacific Air [ATYPICAL SHAPE]

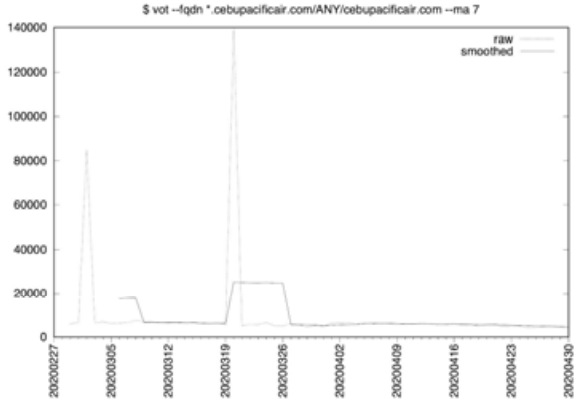

19. China Airlines [ATYPICAL SHAPE]

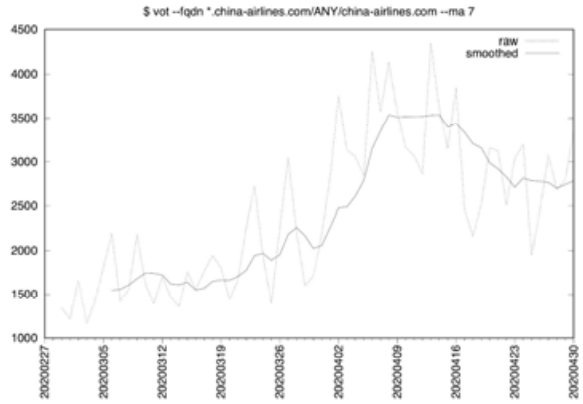

20. Copa Airlines [ATYPICAL SHAPE]

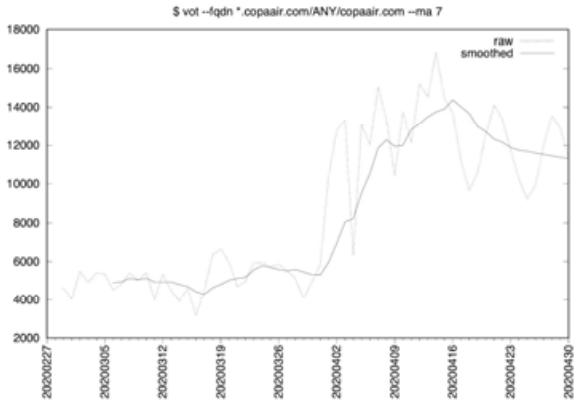

21. China Southern Airlines [ATYPICAL SHAPE]

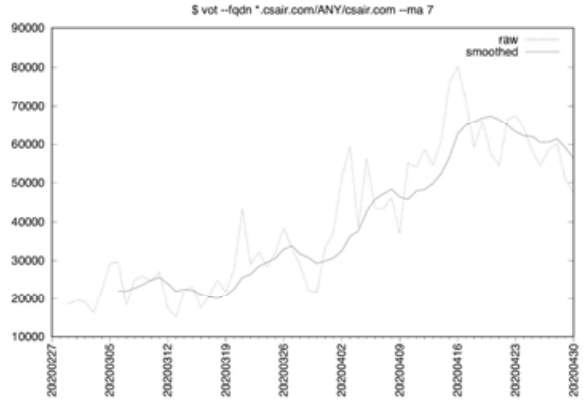

22. Delta Airlines

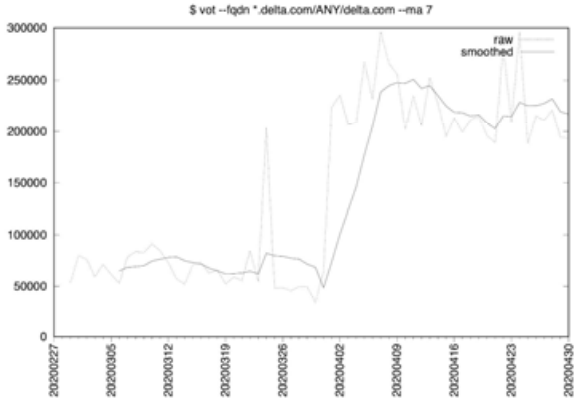

23. EasyJet [ATYPICAL SHAPE]

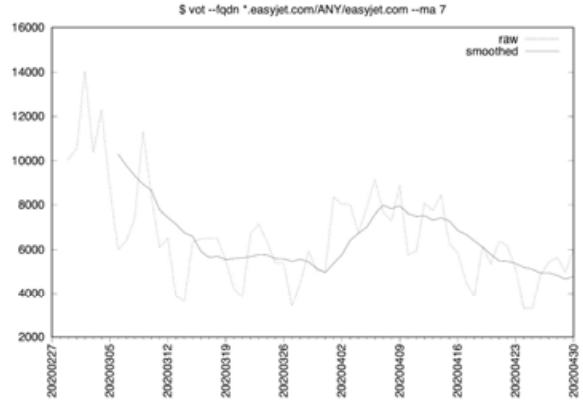

24. EgyptAir [ATYPICAL SHAPE]

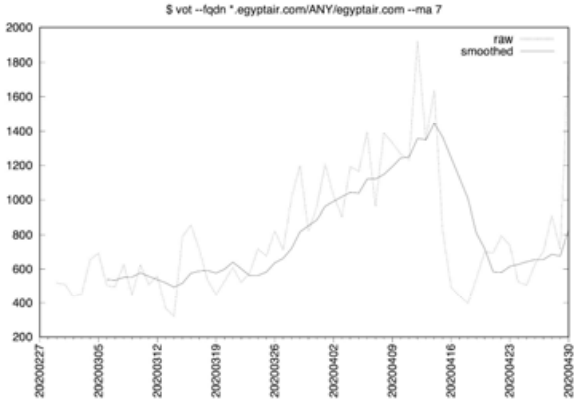

25. El Al [ATYPICAL SHAPE]

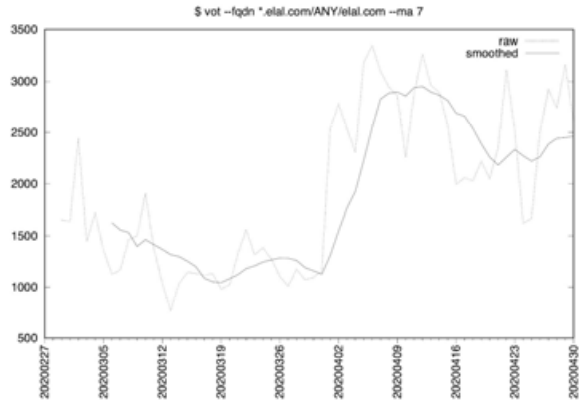

26. Emirates [ATYPICAL SHAPE]

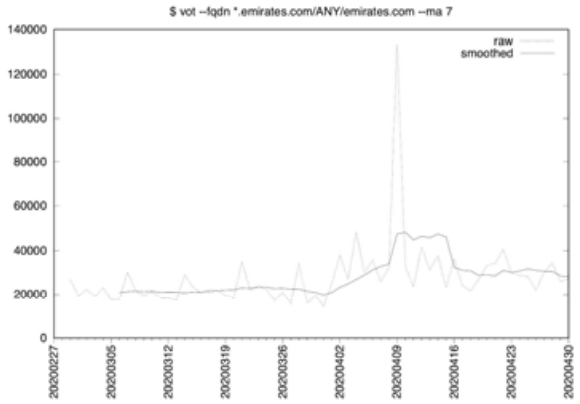

27. Ethioian Airlines [ATYPICAL SHAPE]

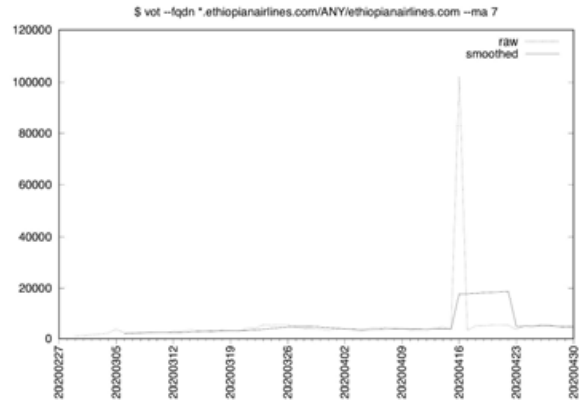

28. Etihad [ATYPICAL SHAPE]

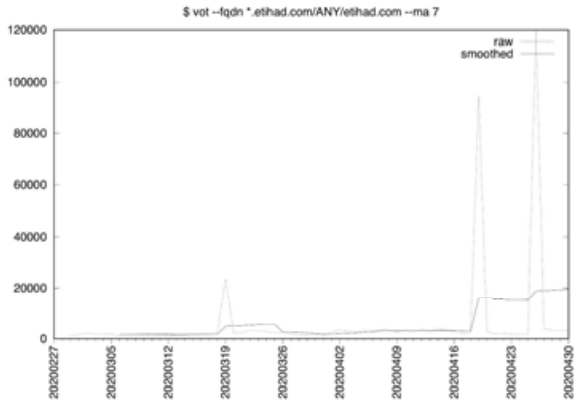

29. Eurowings [ATYPICAL SHAPE]

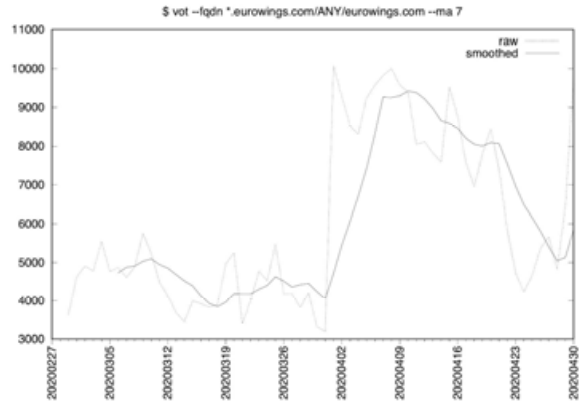

30. EVA Air [ATYPICAL SHAPE]

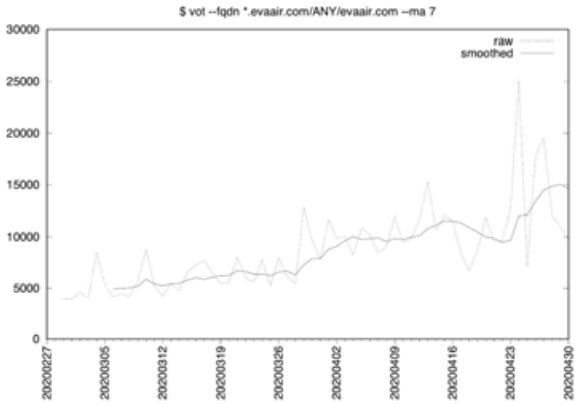

31. Asiana Airlines [ATYPICAL SHAPE]

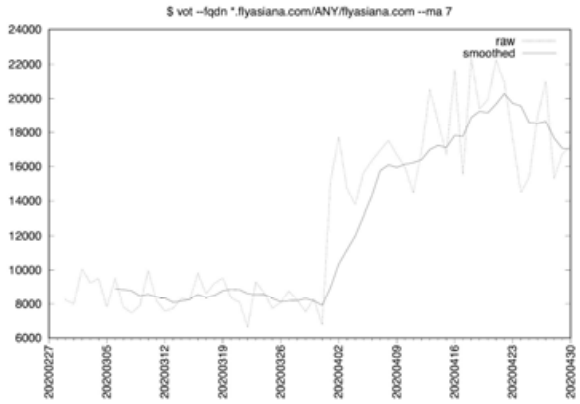

32. Frontier Airlines [ATYPICAL SHAPE]

33. LEVEL [ATYPICAL SHAPE]

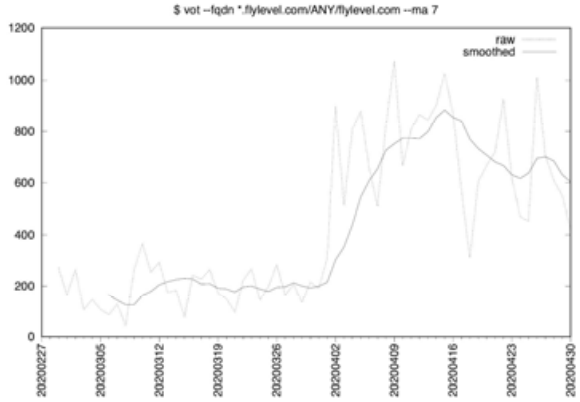

34. Scandanavian Airlines

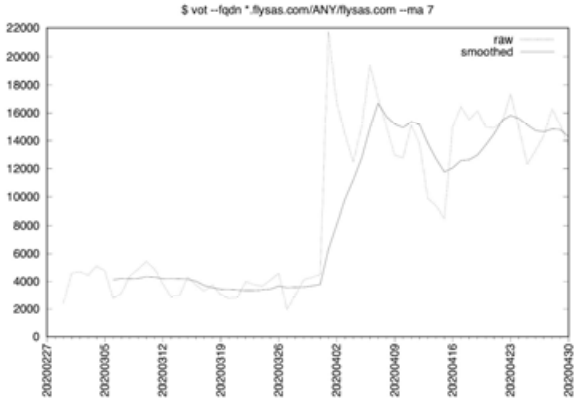

35. TAP Air Portugal [ATYPICAL SHAPE]

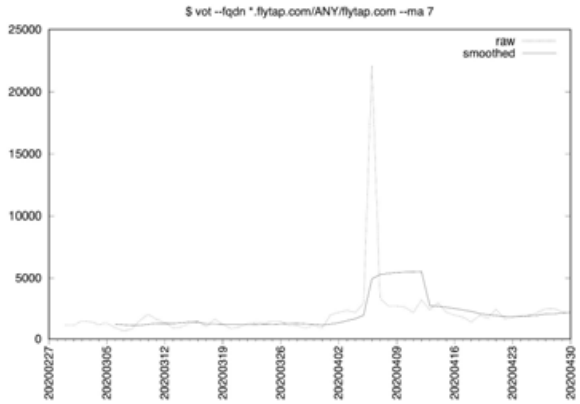

36. Garuda Indonesian Airlines [ATYPICAL SHAPE]

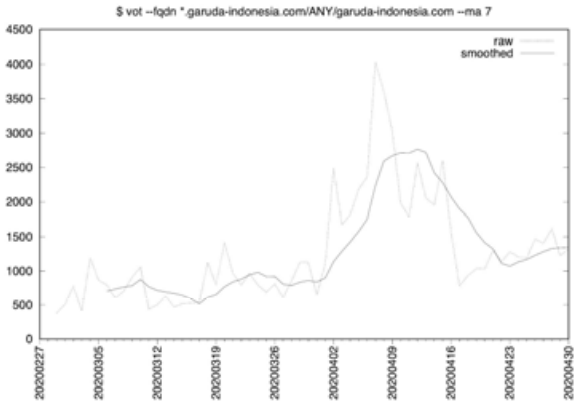

37. Go Airlines (India) [ATYPICAL SHAPE]

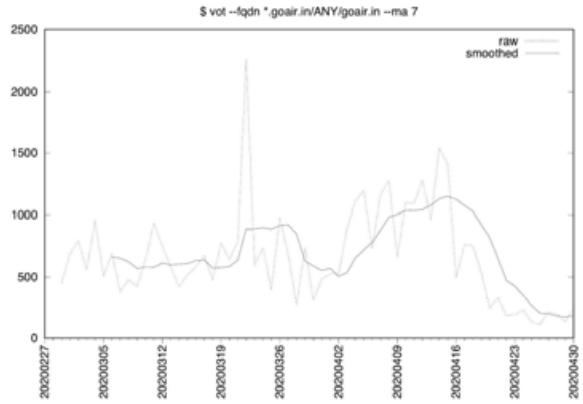

38. IndiGo (India) [ATYPICAL SHAPE]

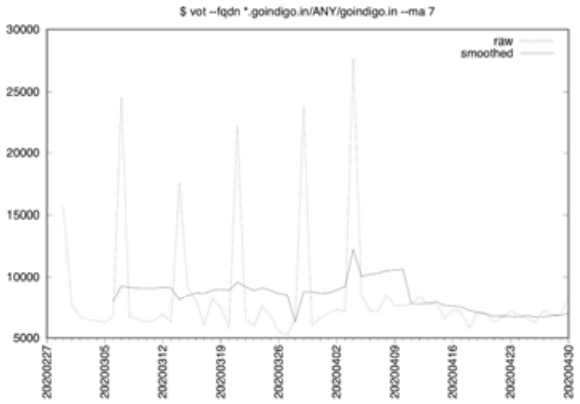

39. Hainan Airlines [ATYPICAL SHAPE]

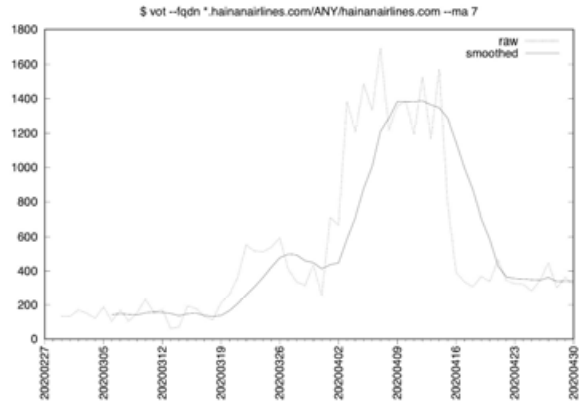

40. Hawaiian Airlines

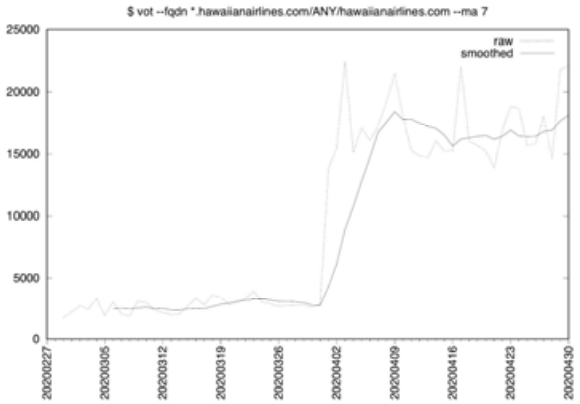

41. Iberia Airlines [ATYPICAL SHAPE]

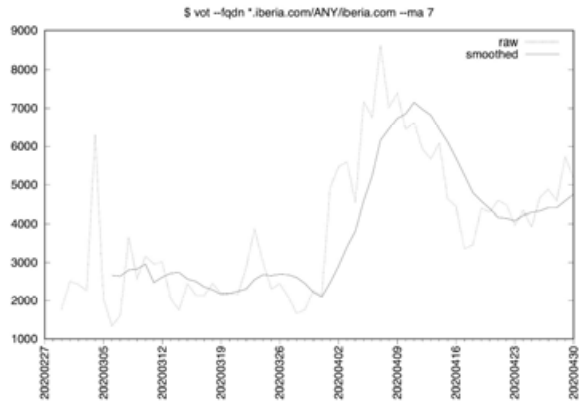

42. Icelandair [ATYPICAL SHAPE]

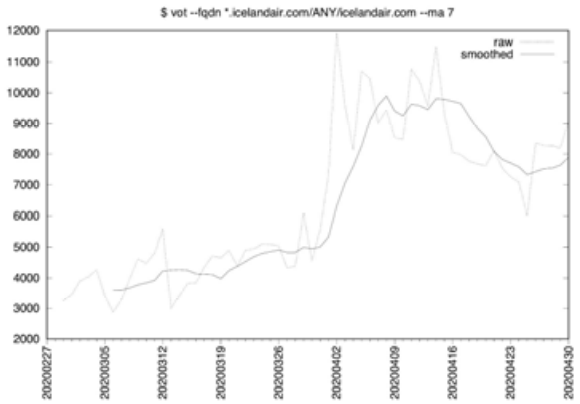

43. Iran Air [ATYPICAL SHAPE]

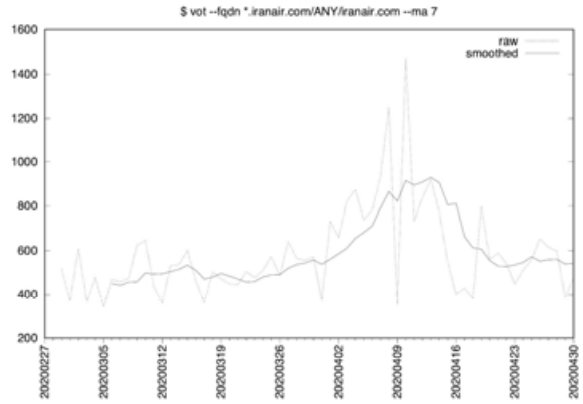

44. Japan Airlines [ATYPICAL SHAPE]

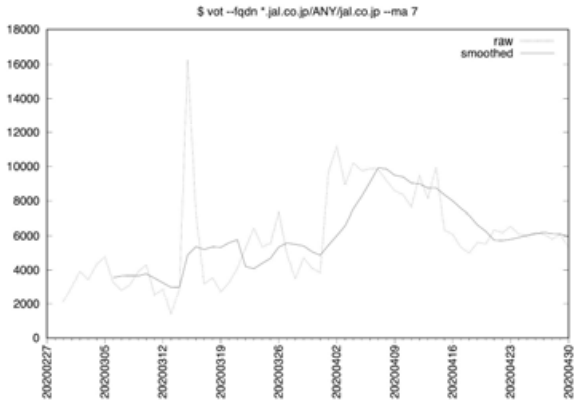

45. JetBlue

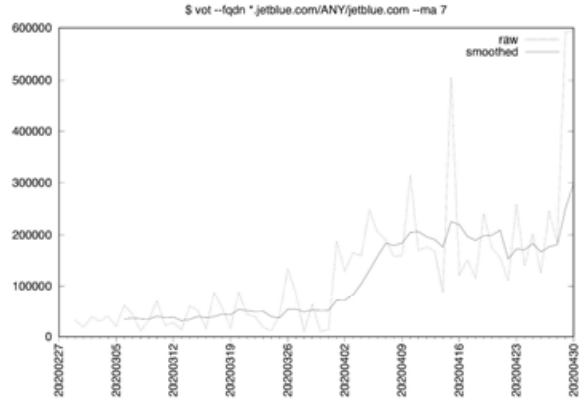

46. KLM

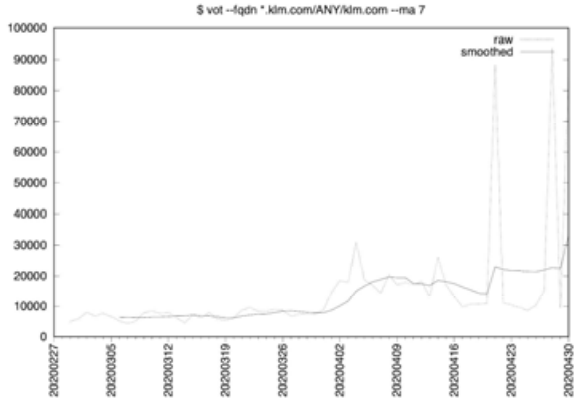

47. Korea Air [ATYPICAL SHAPE]

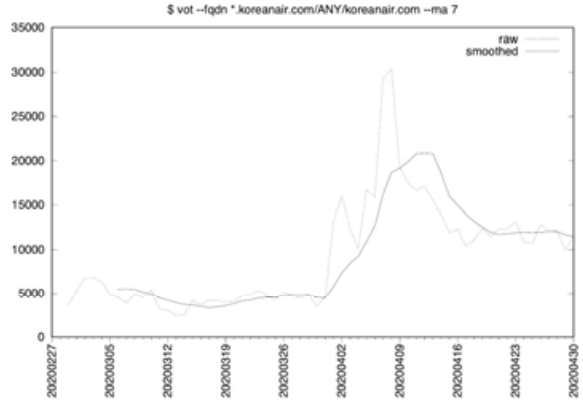

48. LATAM Airlines [ATYPICAL SHAPE]

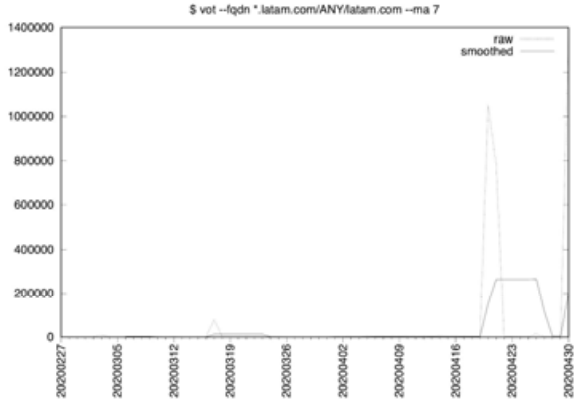

49. Lion Air [ATYPICAL SHAPE]

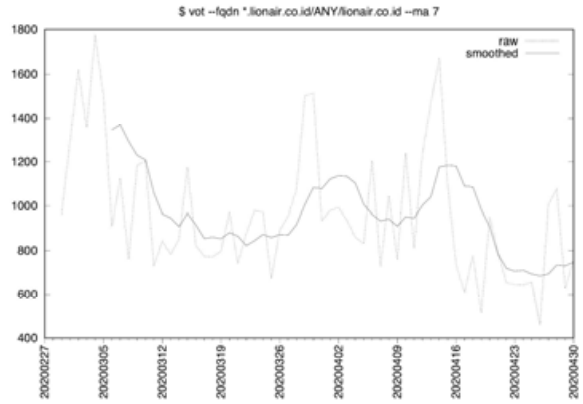

50. Lufthansa [ATYPICAL SHAPE]

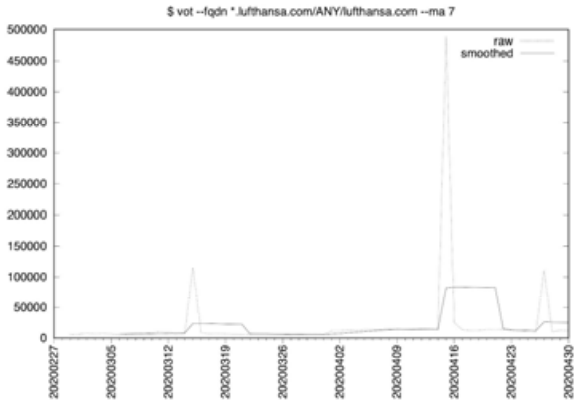

51. Malaysia Airlines [ATYPICAL SHAPE]

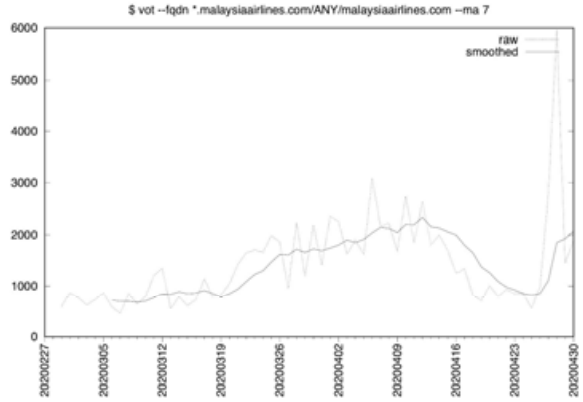

52. Norwegian Airlines

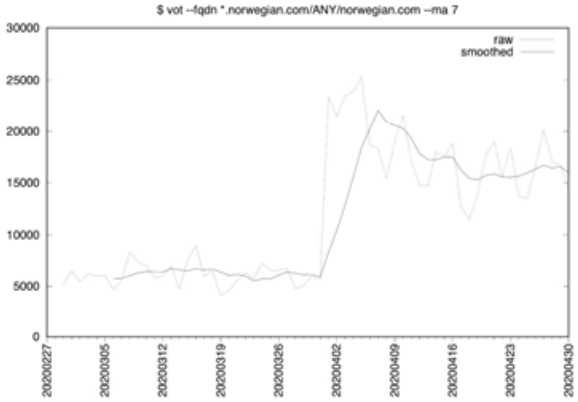

53. Philippine Airline

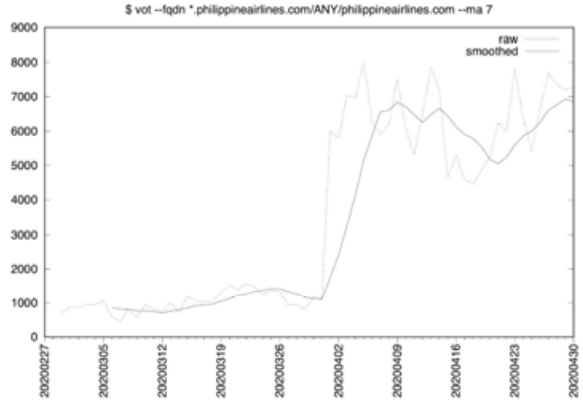

54. Pobeda Airlines [ATYPICAL SHAPE]

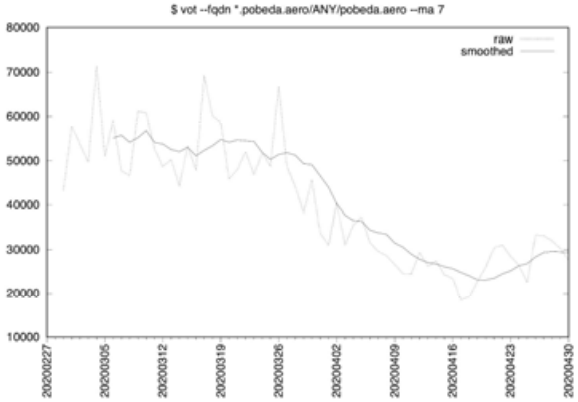

55. Qantas [ATYPICAL SHAPE]

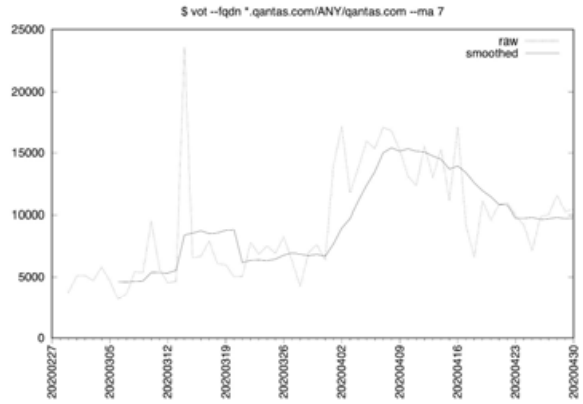

56. Qatar Airways [ATYPICAL SHAPE]

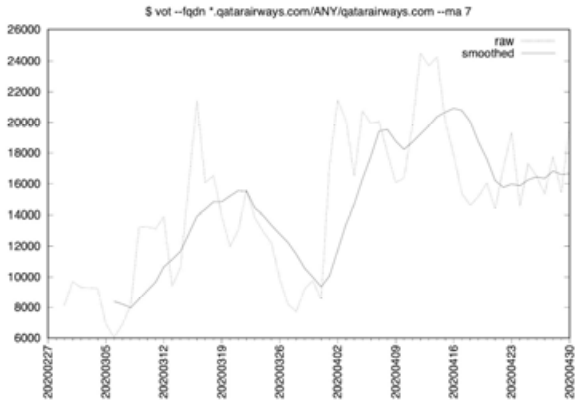

57. Rossiya Airlines [ATYPICAL SHAPE]

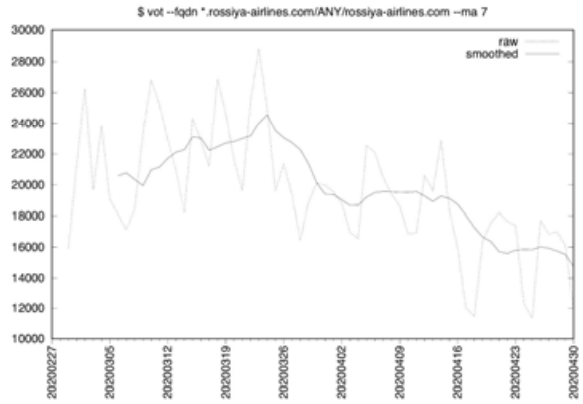

58. RyanAir [ATYPICAL SHAPE]

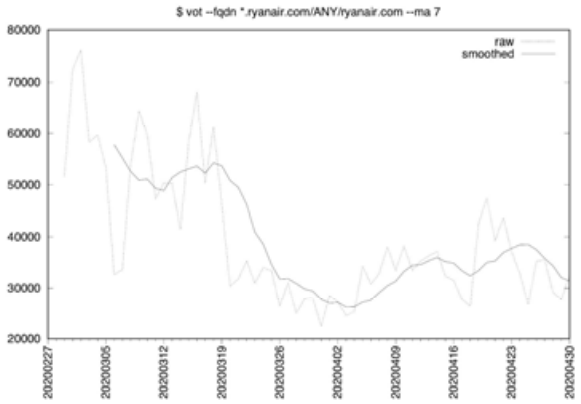

59. S7 Airlines (JSC Siberia Airlines) [ATYPICAL SHAPE]

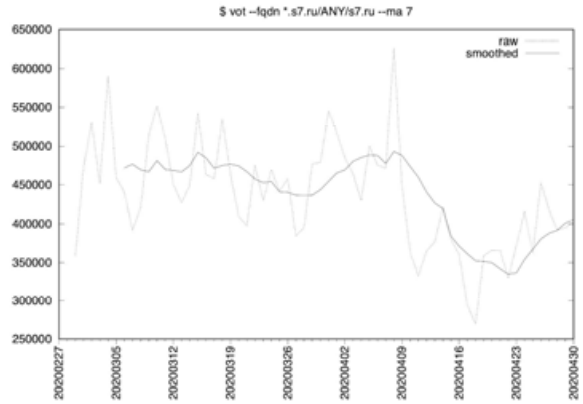

60. Saudia [ATYPICAL SHAPE]

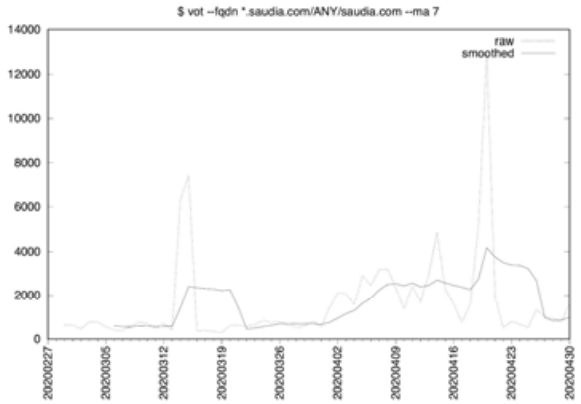

61. Singapore Airlines

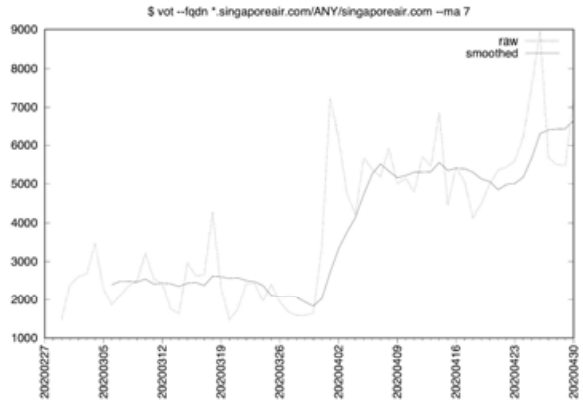

62. Southwest

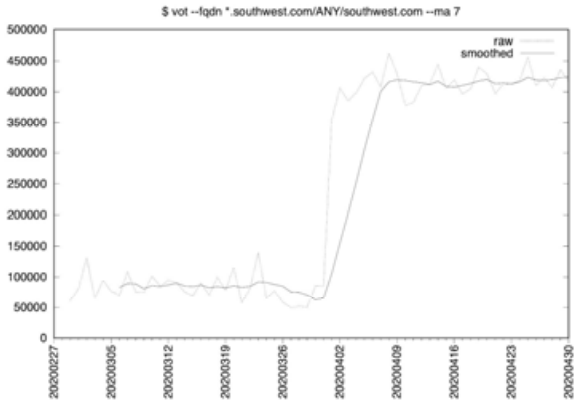

63. SpiceJet (India) [ATYPICAL SHAPE]

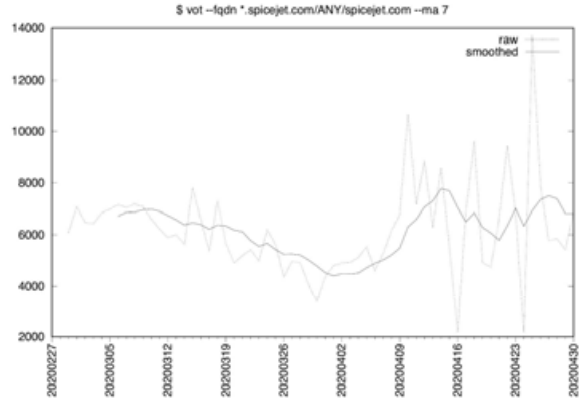

64. Spirit Airlines

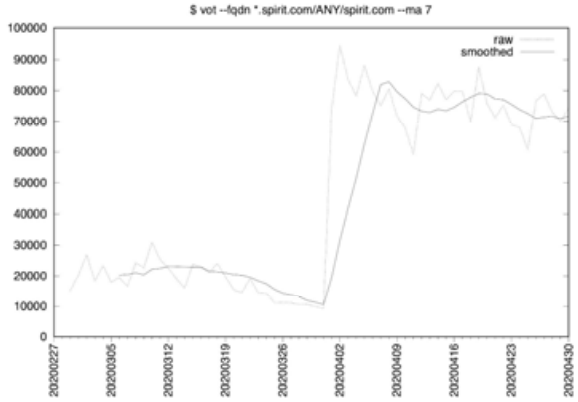

65. Swissair [ATYPICAL SHAPE]

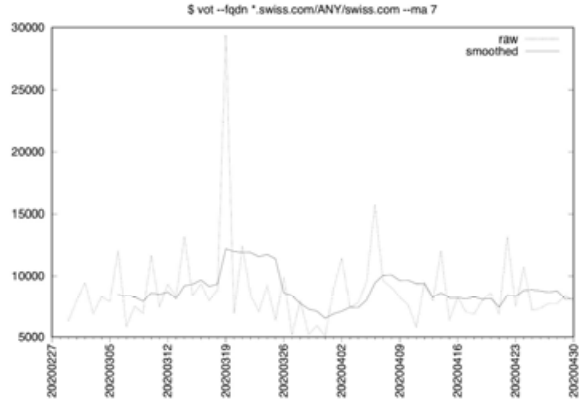

66. Thai Airways [ATYPICAL SHAPE]

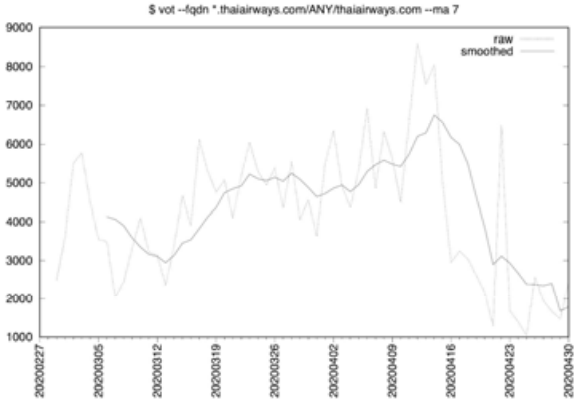

67. Transavia Airlines [ATYPICAL SHAPE]

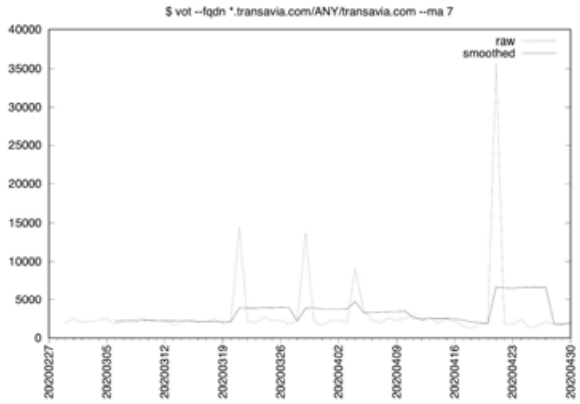

68. Turkish Airlines [ATYPICAL SHAPE]

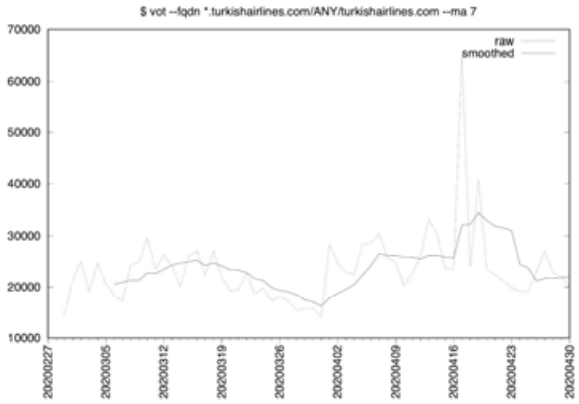

69. United Airlines [ATYPICAL SHAPE]

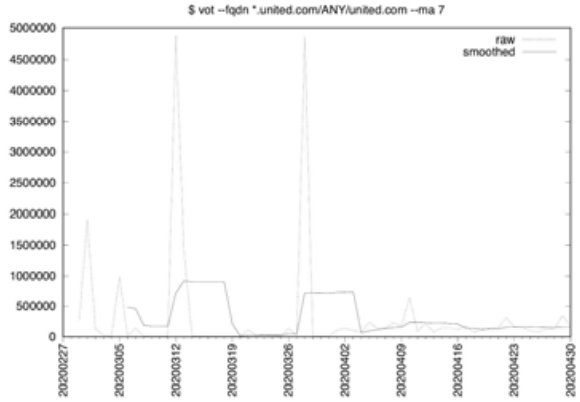

70. VietJet Air [ATYPICAL SHAPE]

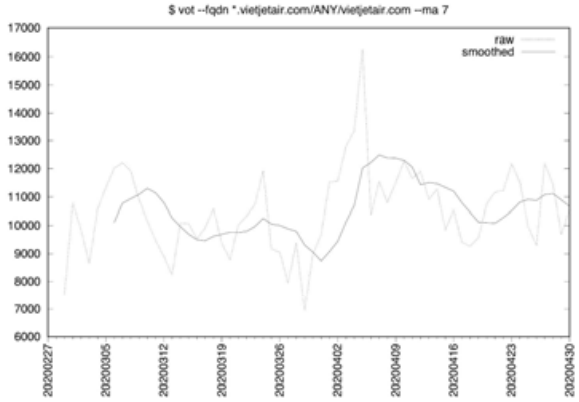

71. Vietnam Airlines [ATYPICAL SHAPE]

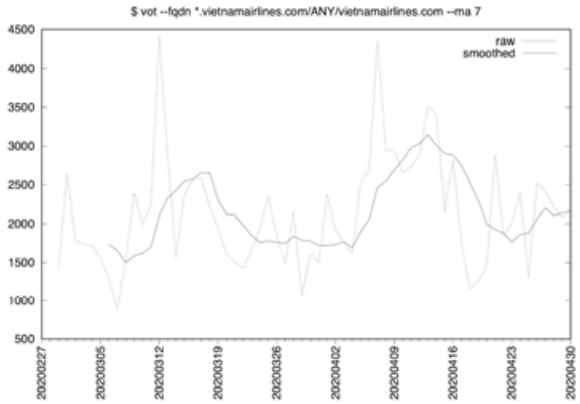

72. Vueling Airlines [ATYPICAL SHAPE]

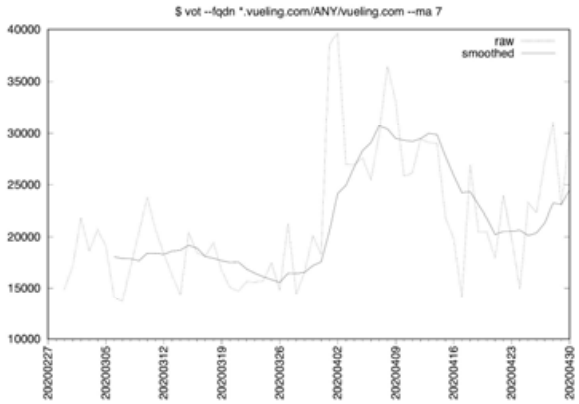

73. WizzAir [ATYPICAL SHAPE]

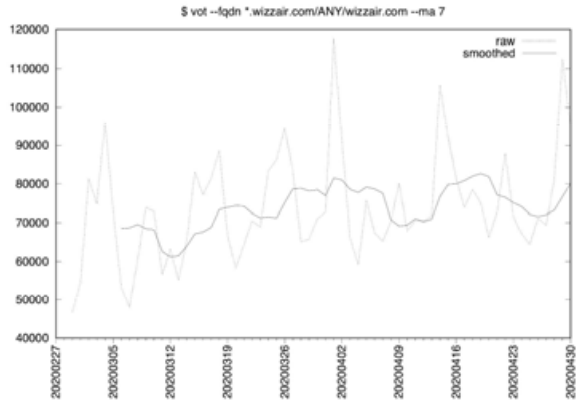

Car Rental Companies

1. alamo.com
2. avis.com
3. budget.com
4. dollar.com
5. enterprise.com
6. hertz.com
7. nationalcar.com

1. Alamo

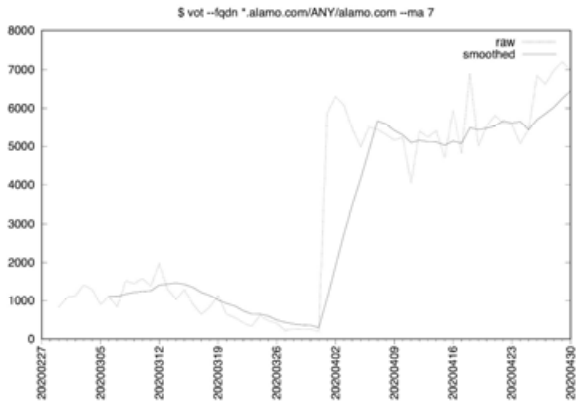

2. Avis

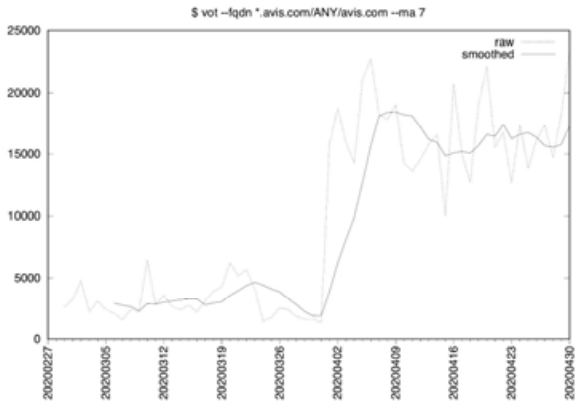

Budget

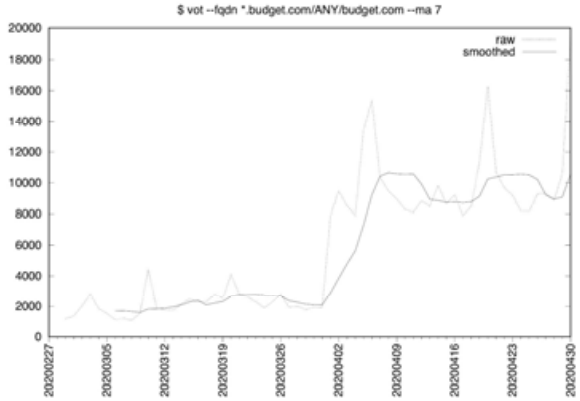

4. Dollar

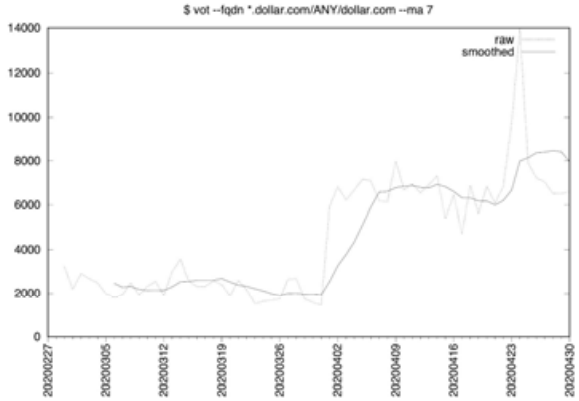

5. Enterprise

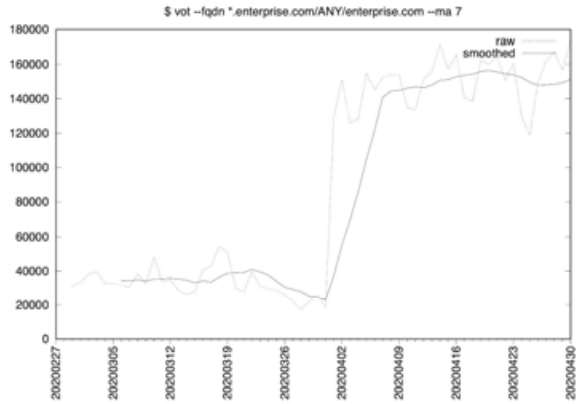

6. Hertz

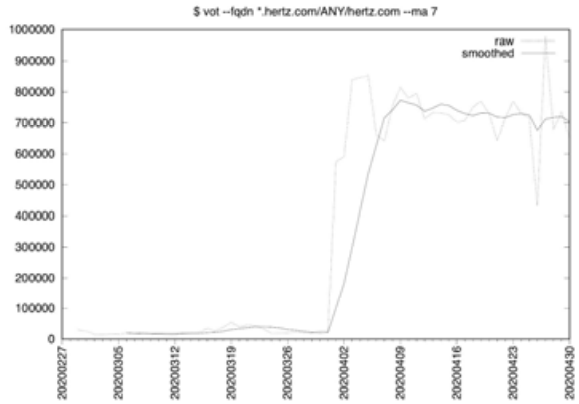

7. National Car Rental [ATYPICAL SHAPE]

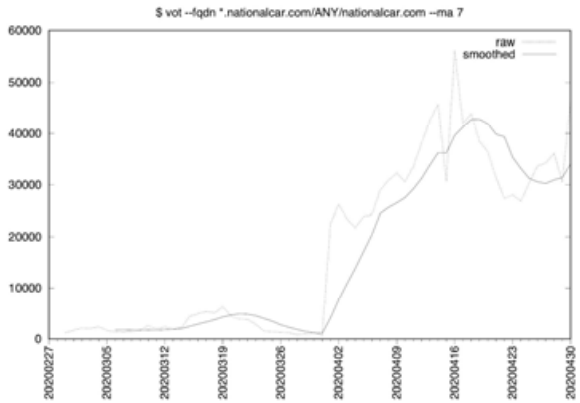

Cruise Lines

1. carnival.com
2. hollandamerica.com
3. ncl.com
4. princess.com
5. royalcaribbean.com

1. Carnival Cruise Lines [ATYPICAL SHAPE]

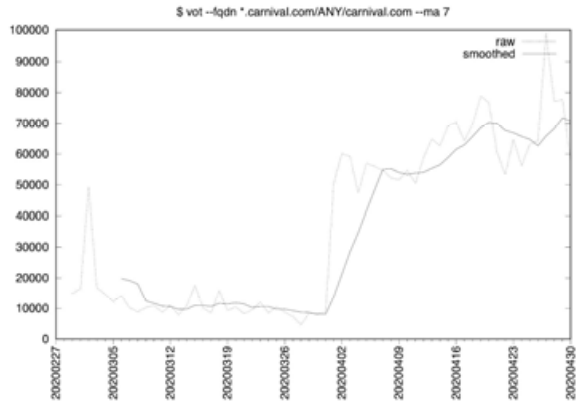

2. Holland America Line

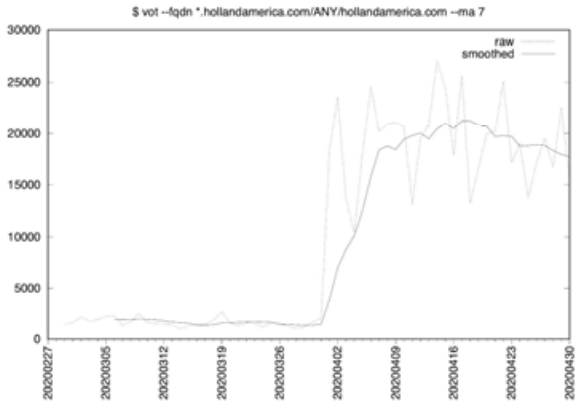

3. Norwegian Cruise Lines

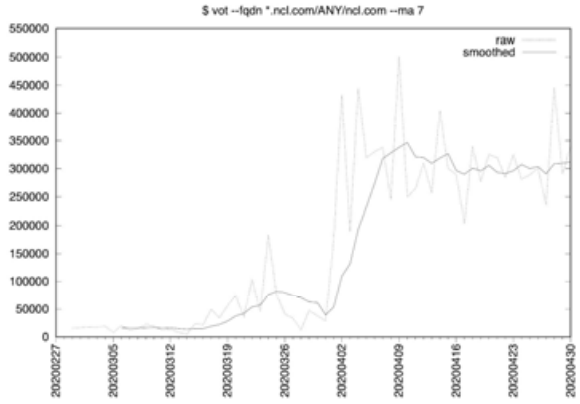

4. Princess Cruise Lines

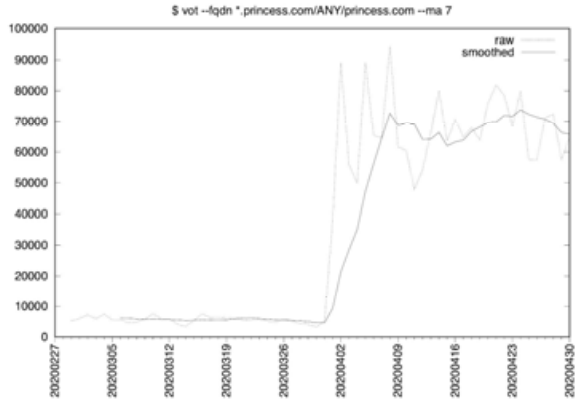

5. Royal Caribbean Cruise Lines

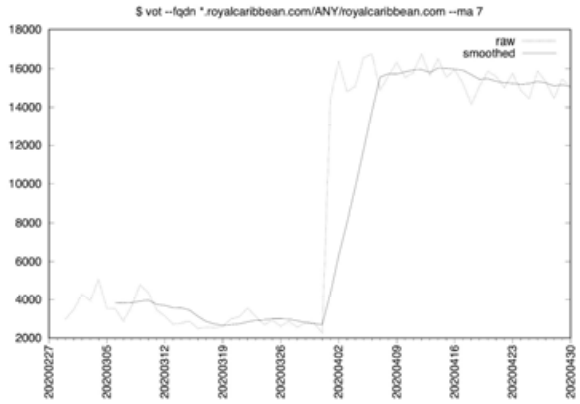

Hotels

1. accor.com
2. bestwestern.com
3. choicehotels.com
4. hilton.com
5. hyatt.com
6. ihg.com
7. marriott.com
8. wyndhamhotels.com

1. Accor [ATYPICAL SHAPE]

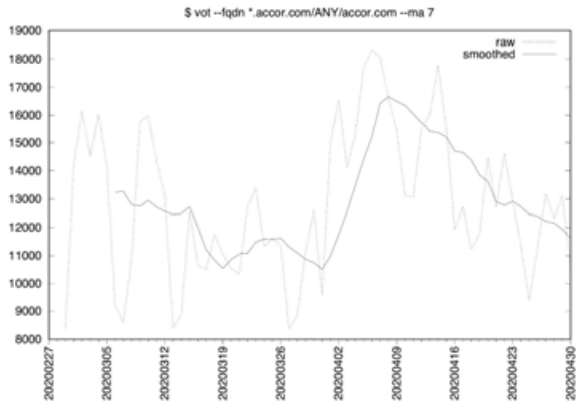

2. Best Western

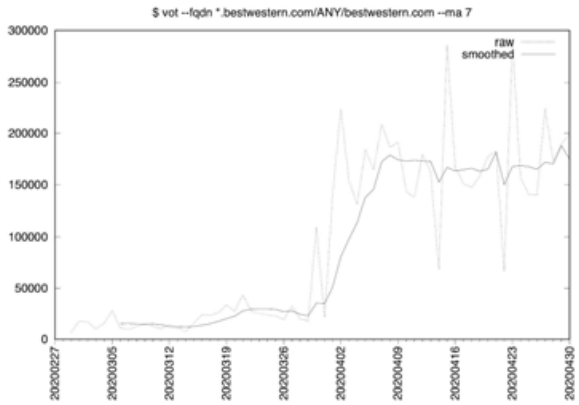

3. Choice Hotels

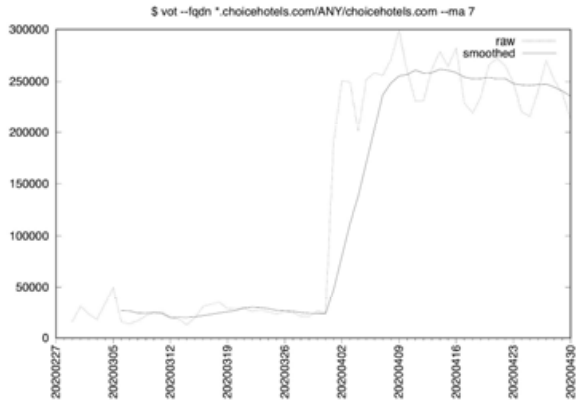

4. Hilton

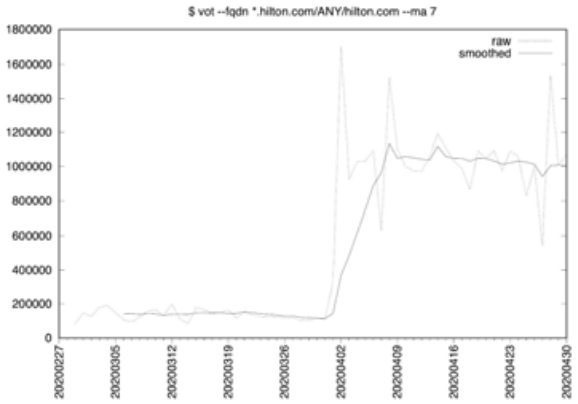

5. Hyatt [ATYPICAL SHAPE]

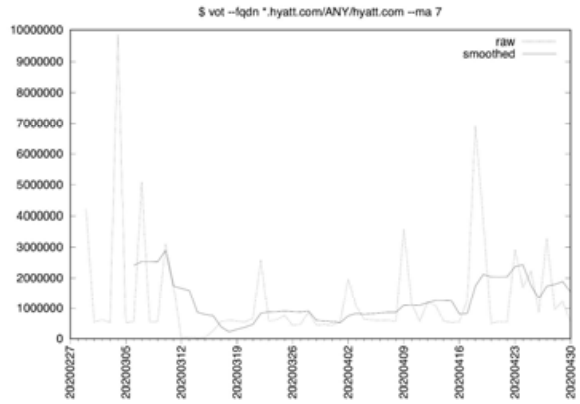

6. IHG

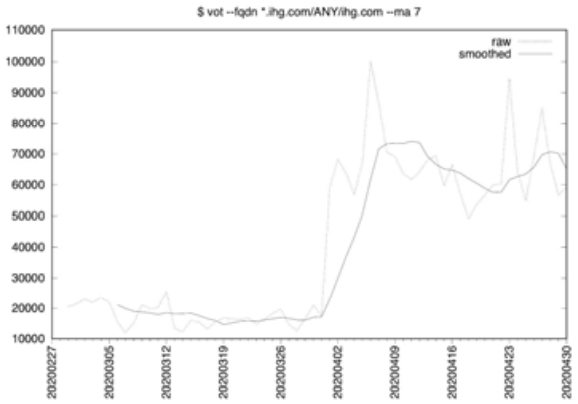

7. Marriott

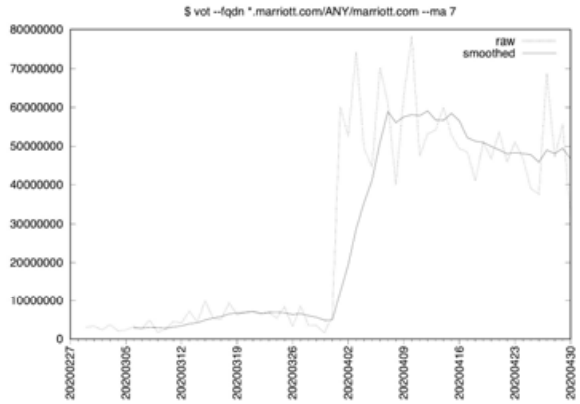

8. Wyndham Hotels

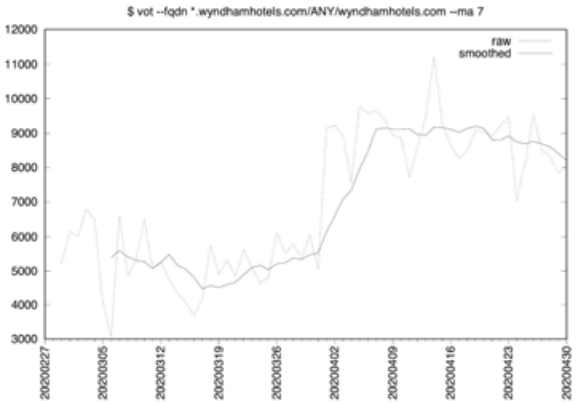

Railroads

1. amtrak.com
2. bnsf.com
3. csx.com
4. kcsouthern.com
5. nscorp.com
6. up.com

1. Amtrak

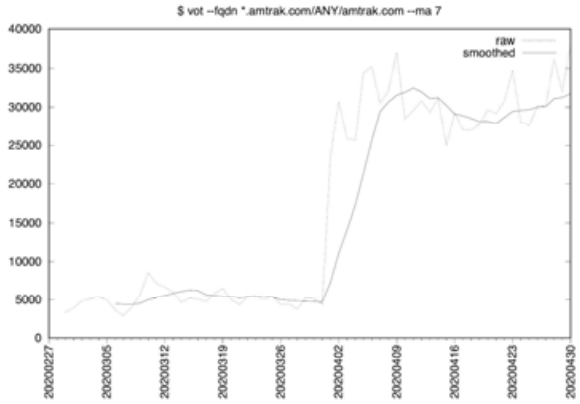

2. Burlington Northern

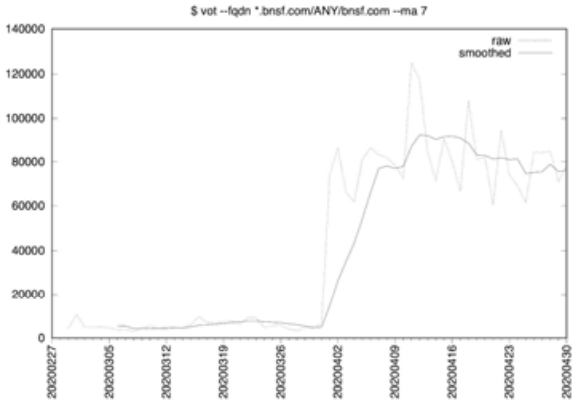

3. CSX [ATYPICAL SHAPE]

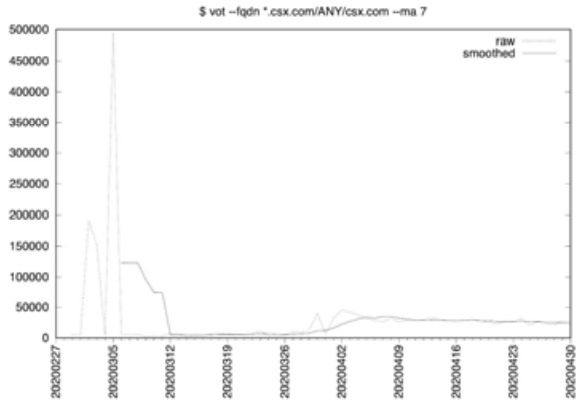

4. Kansas City Southern

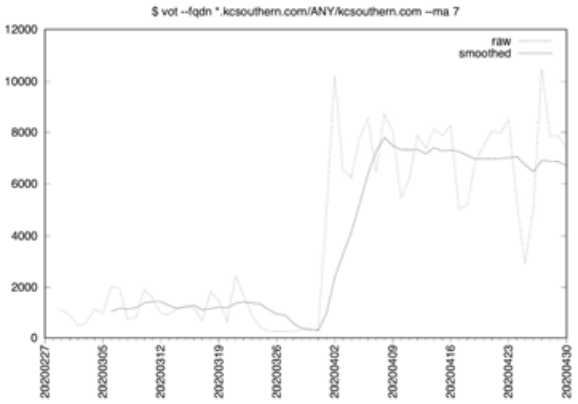

5. Norfolk Southern [ATYPICAL SHAPE]

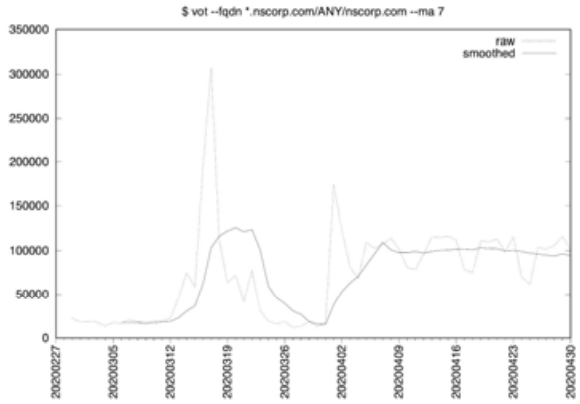

6. Union Pacific

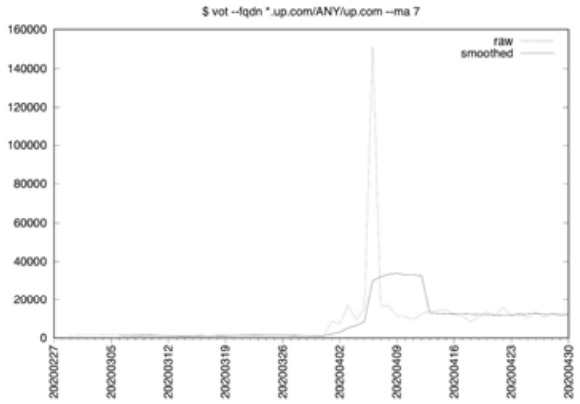

Ride Sharing Companies

1. Lyft

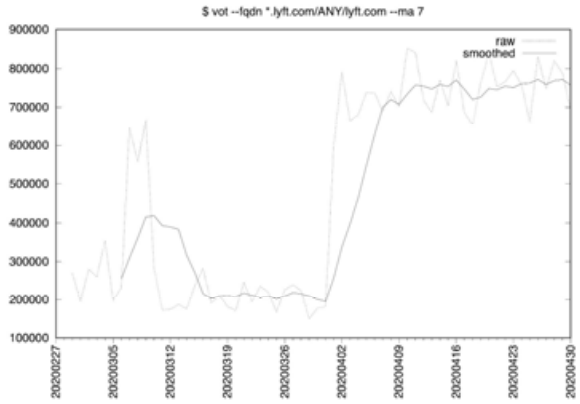

2. Uber [ATYPICAL SHAPE]

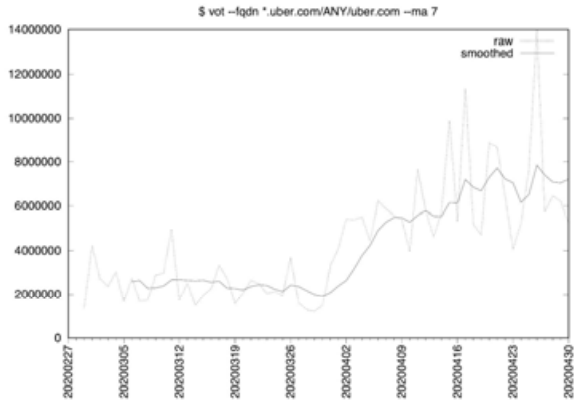

Shipping/Logistics

1. DHL

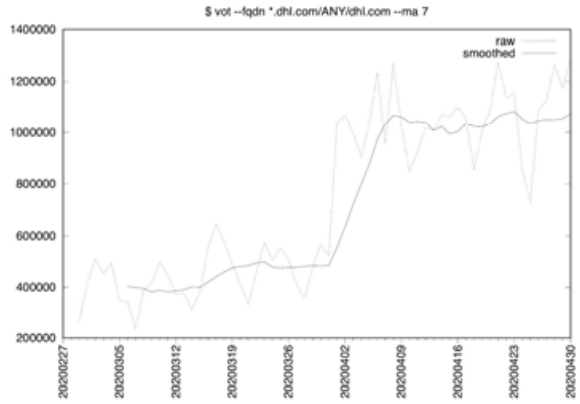

2. Fedex

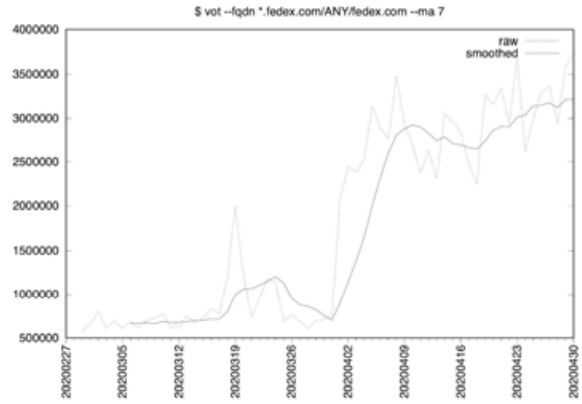

3. UPS [ATYPICAL SHAPE]

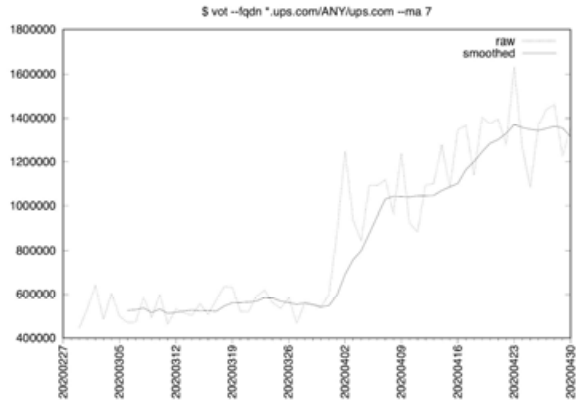

4. USPS

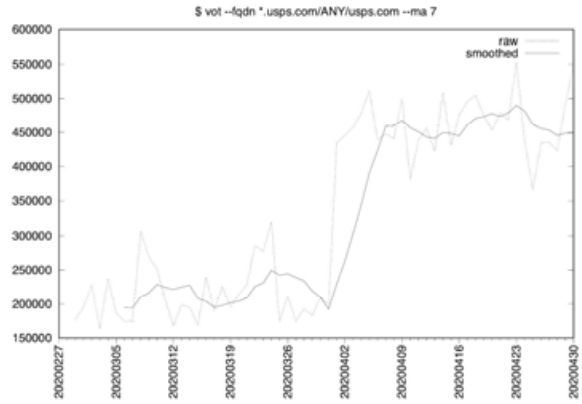

Retail Sites

Apparel and Department Stores

1. bloomingdales.com
2. dillards.com
3. gap.com
4. jcpenney.com
5. kohls.com
6. lordandtaylor.com
7. macys.com
8. nordstrom.com
9. rossstores.com
10. sears.com
11. target.com
12. walmart.com

1. Bloomingdales

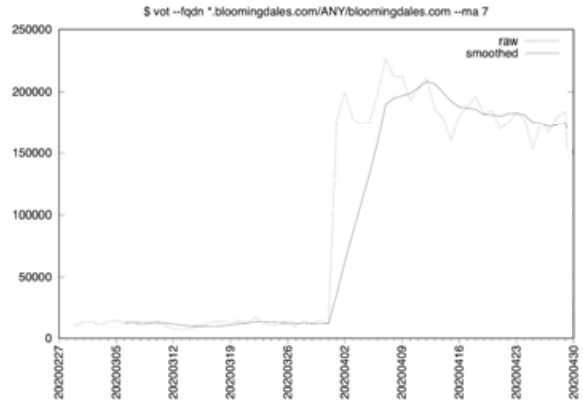

2. Dillards

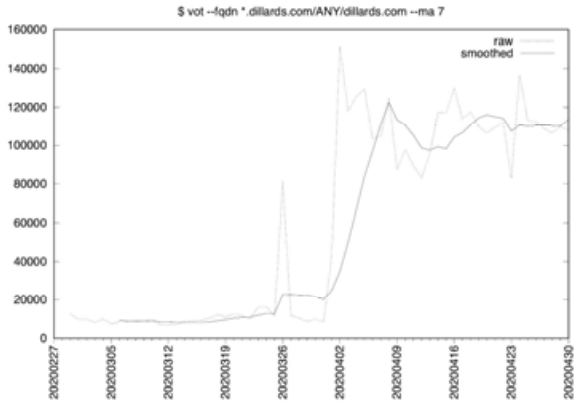

3. Gap

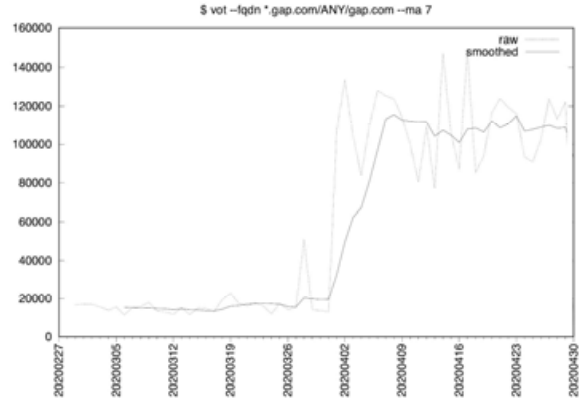

4. J.C. Penney

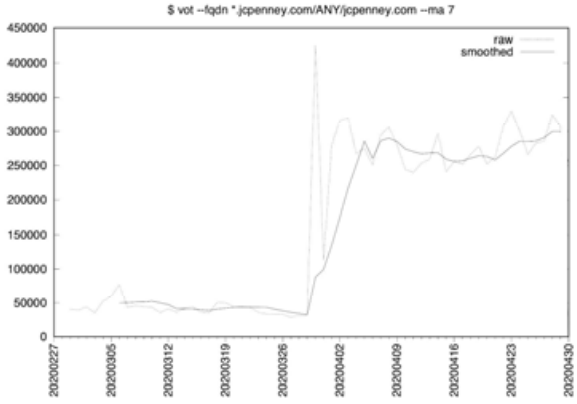

5. Kohls [ATYPICAL SHAPE]

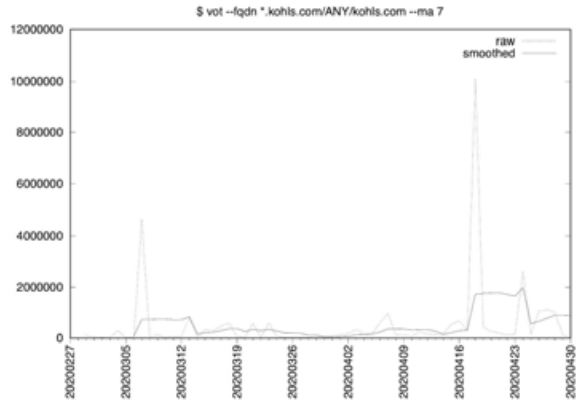

6. Lord and Taylor [ATYPICAL SHAPE]

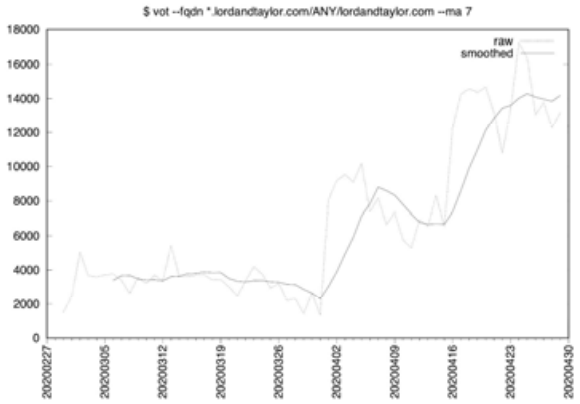

7. Macy's [ATYPICAL SHAPE]

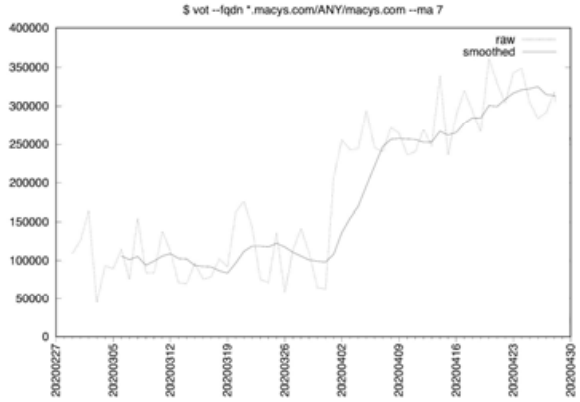

8. Nordstroms

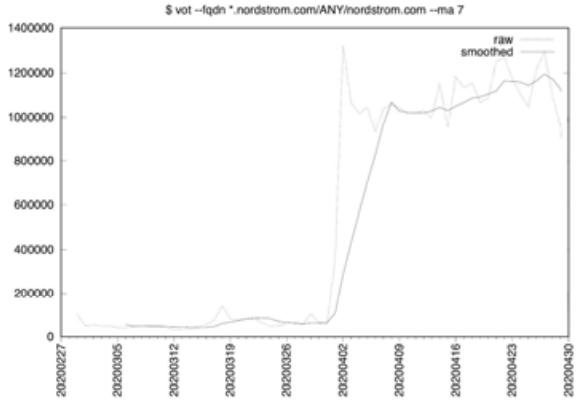

9. Ross Stores

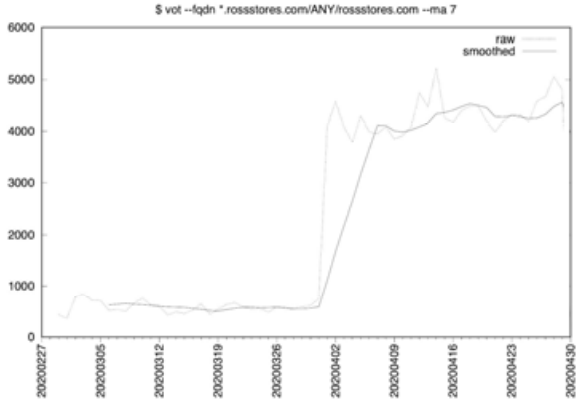

10. Sears

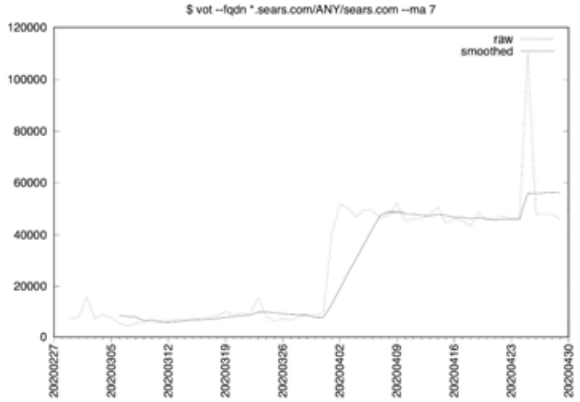

11. Target

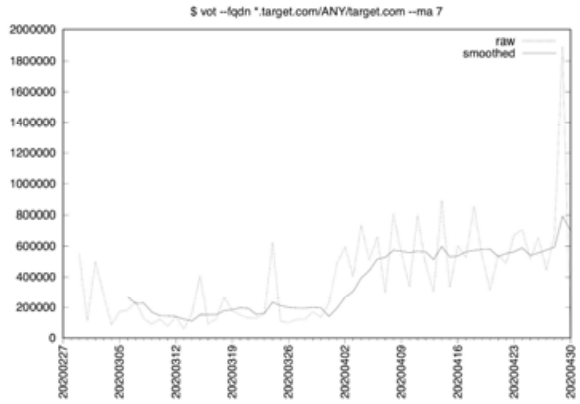

12. Walmart [ATYPICAL SHAPE]

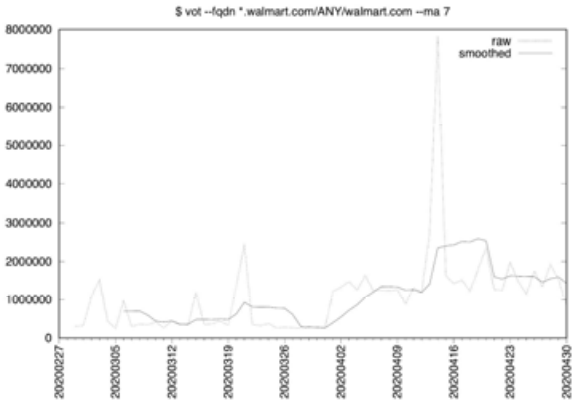

Convenience and Dollar Stores

1. 7-11

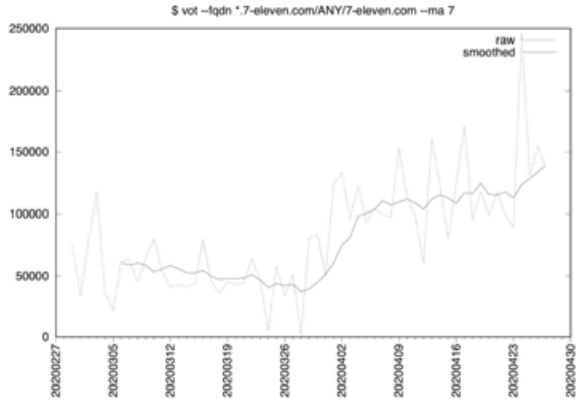

2. Dollar General

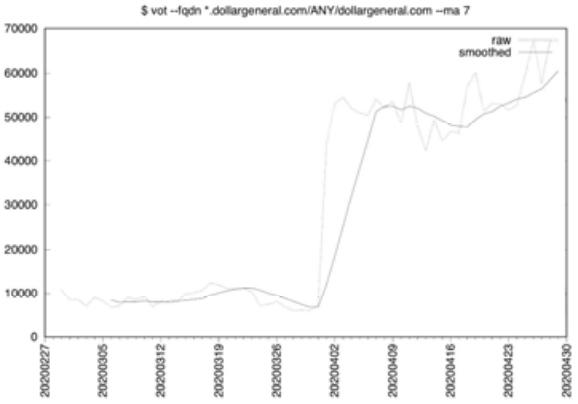

3. Dollar Tree

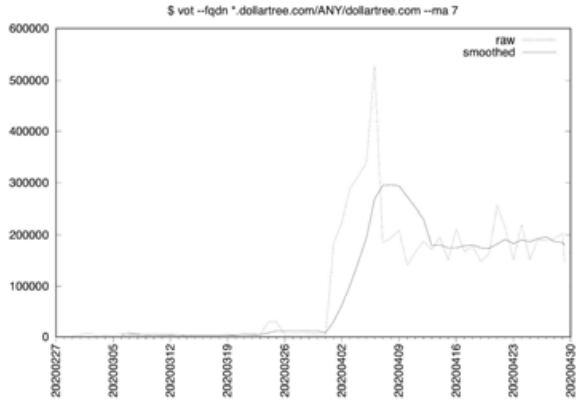

Electronics

1. Apple [ATYPICAL SHAPE]

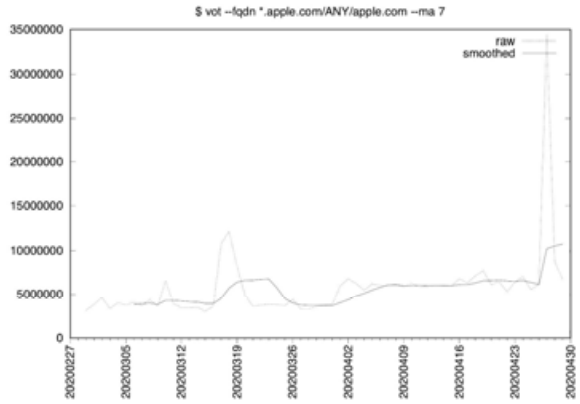

2. Bestbuy [ATYPICAL SHAPE]

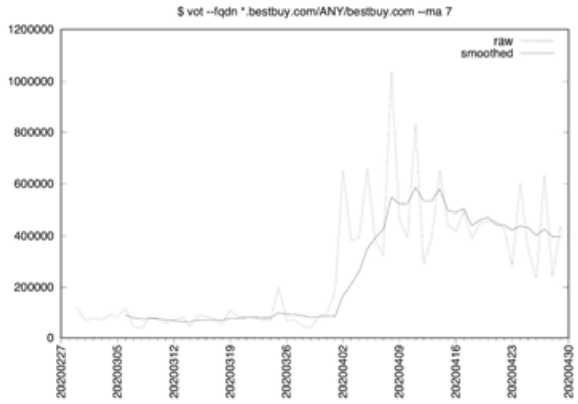

3. Newegg

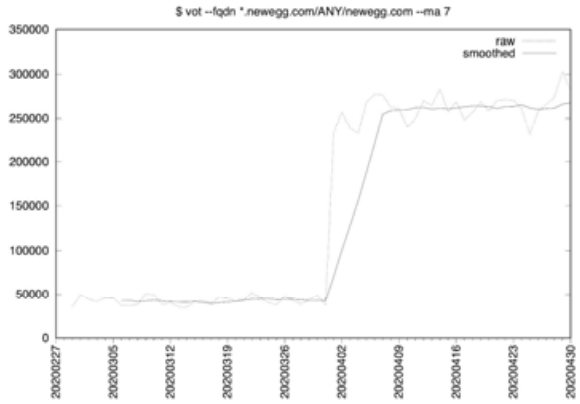

Food Delivery

1. doordash.com

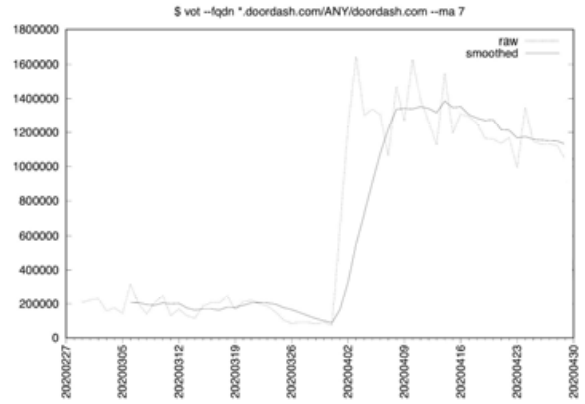

2. grubhub.com

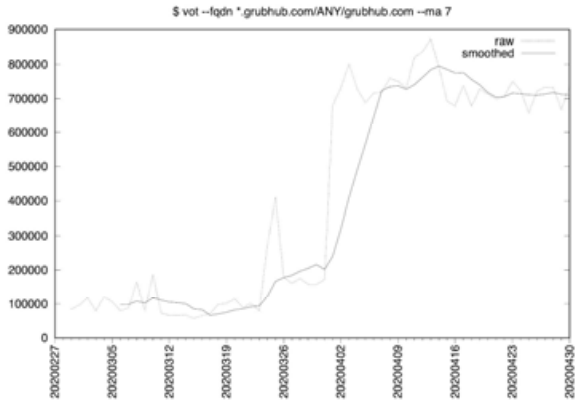

Hardware/Home Improvement/Home Furnishings

1. acehardware.com
2. bedbathandbeyond.com
3. homedepot.com
4. ikea.com
5. lowes.com
6. menards.com
7. truevalue.com
8. wayfair.com

1. Ace Hardware

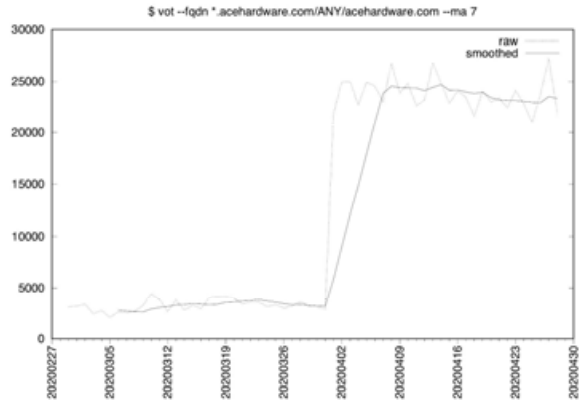

2. Bed Bath and Beyond

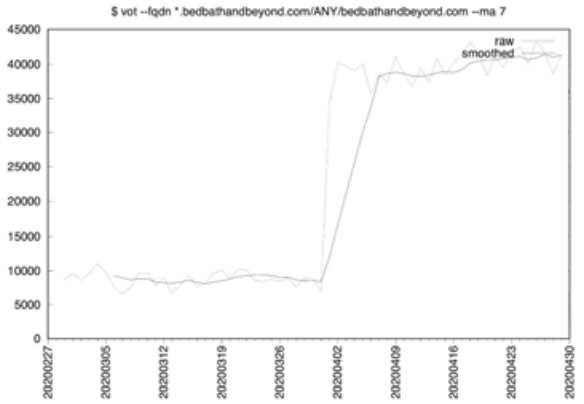

3. Home Depot

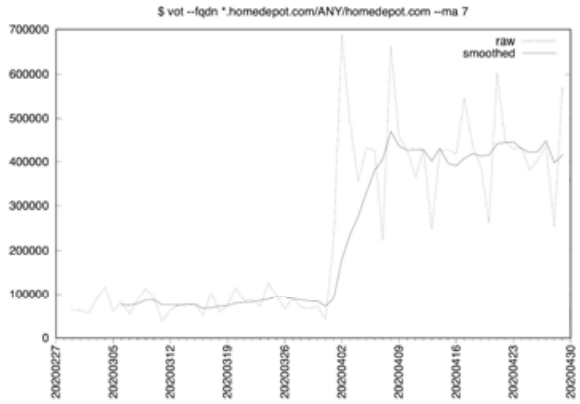

4. Ikea [ATYPICAL SHAPE]

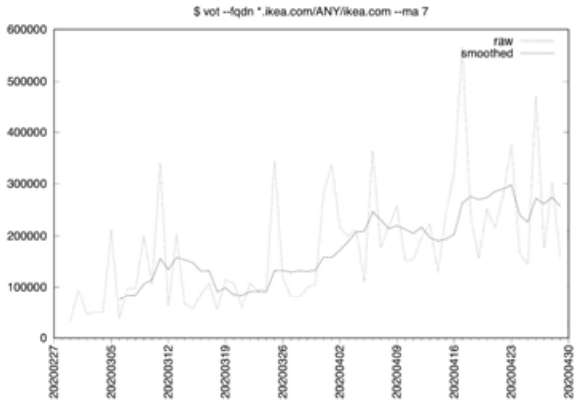

5. Lowes

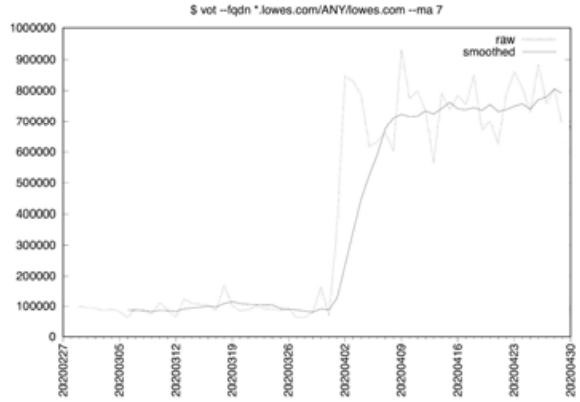

6. Menards

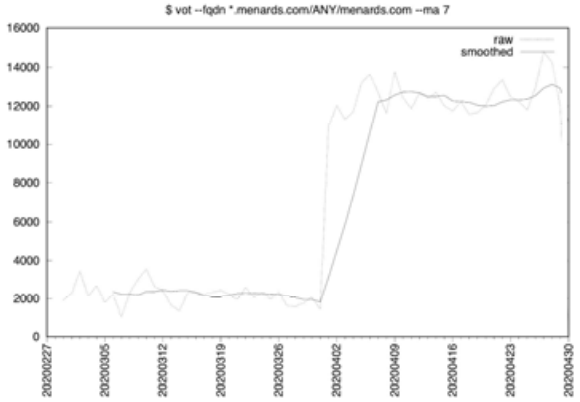

7. True Value

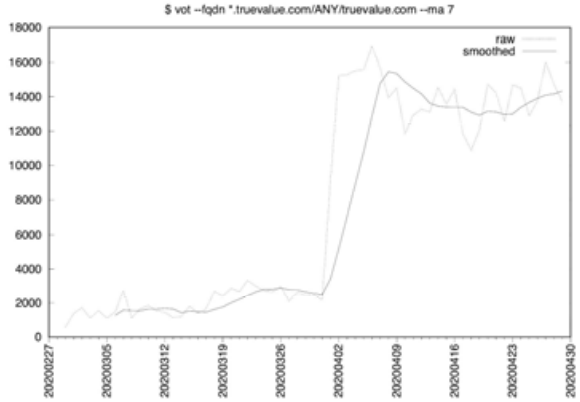

8. Wayfair [ATYPICAL SHAPE]

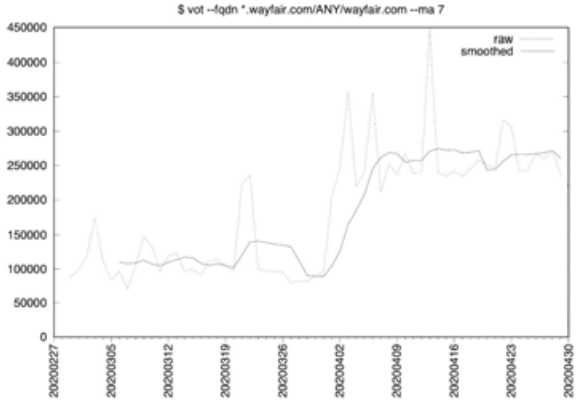

Online Retailers

1. aliexpress.com
2. amazon.com
3. chewy.com
4. ebay.com
5. etsy.com
6. overstock.com
7. qvc.com
8. wish.com

1. Aliexpress [ATYPICAL SHAPE]

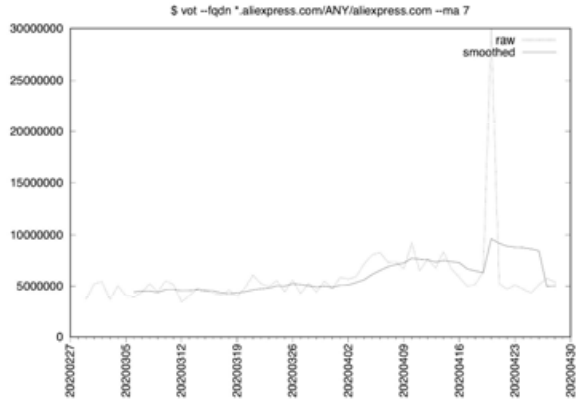

2. Amazon

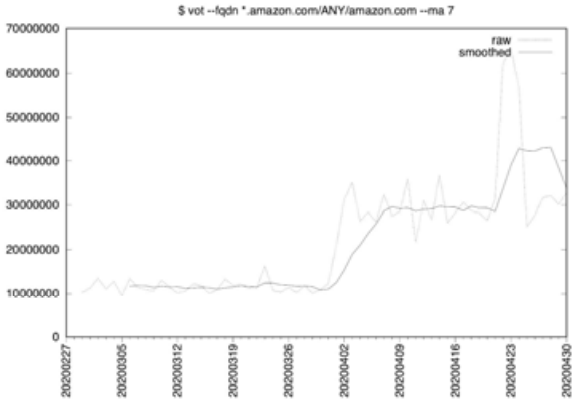

3. Chewy

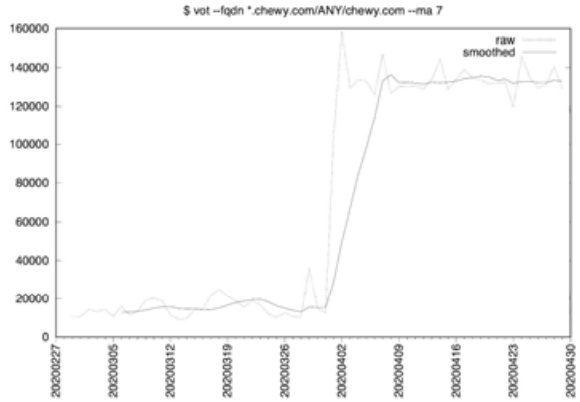

4. Ebay [ATYPICAL SHAPE]

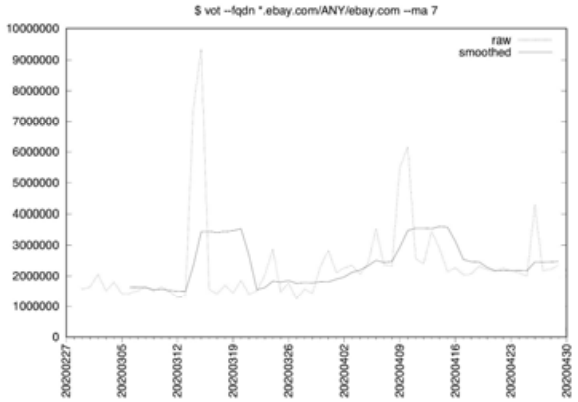

5. Etsy [ATYPICAL SHAPE]

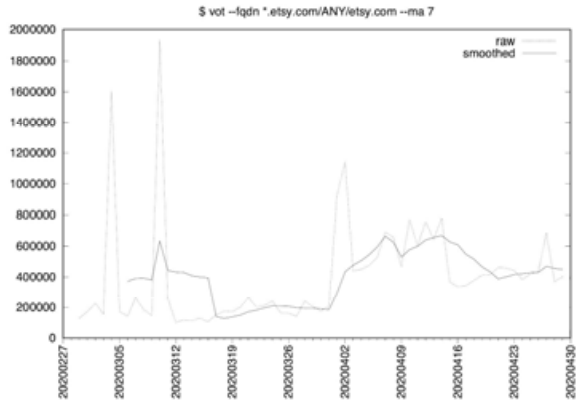

6. Overstock [ATYPICAL SHAPE]

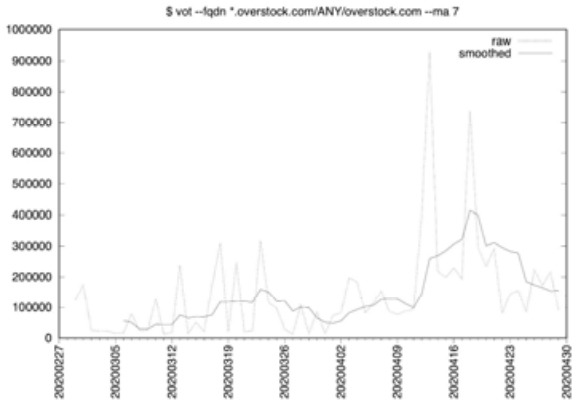

7. QVC

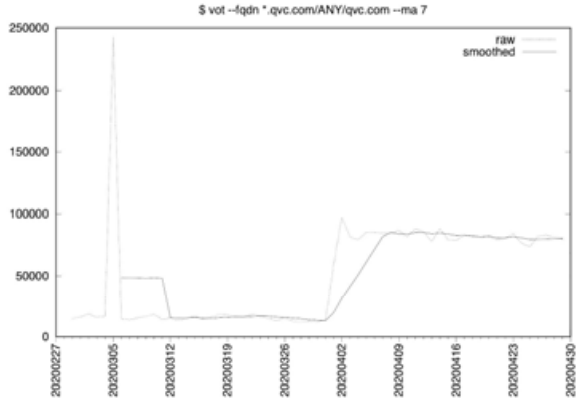

8. Wish

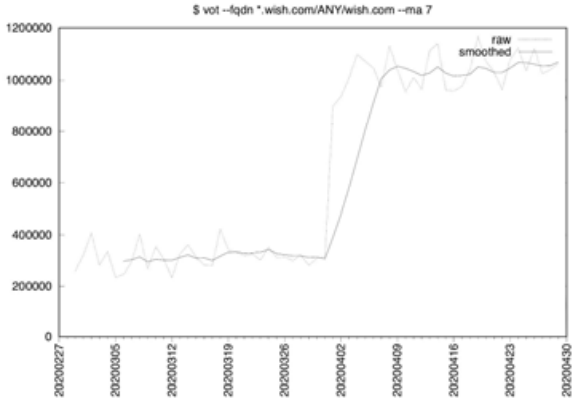

Pharmacy

1. CVS

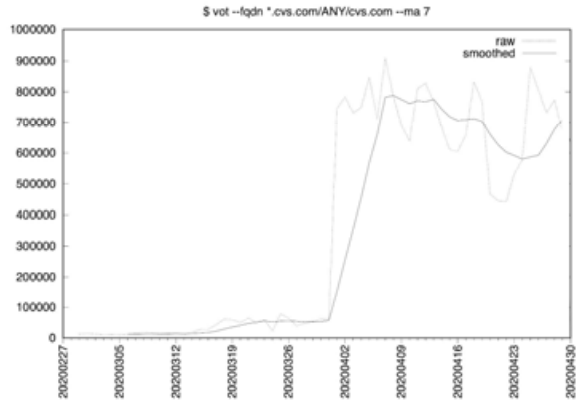

2. Riteaid [ATYPICAL SHAPE]

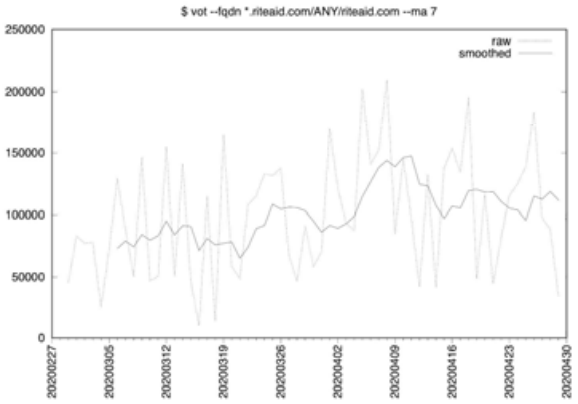

3. Shoprite [ATYPICAL SHAPE]

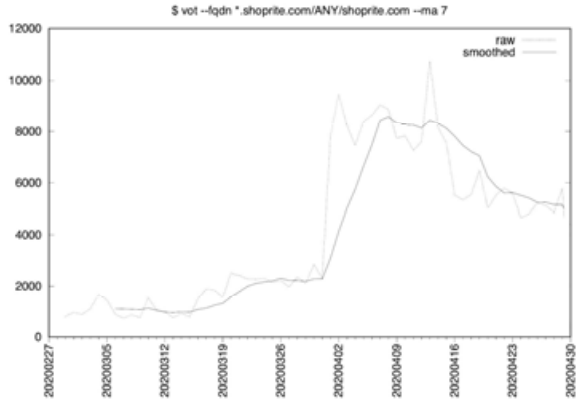

4. Walgreens [ATYPICAL SHAPE]

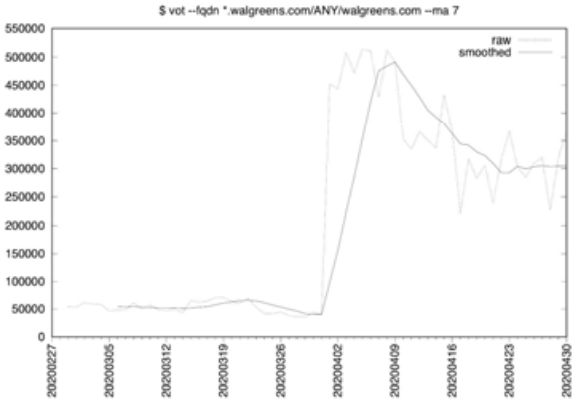

Restaurant Chains

1. arbys.com
2. bk.com
3. chick-fil-a.com
4. mcdonalds.com
5. starbucks.com
6. subway.com
7. tacobell.com
8. wendys.com

1. Arbys

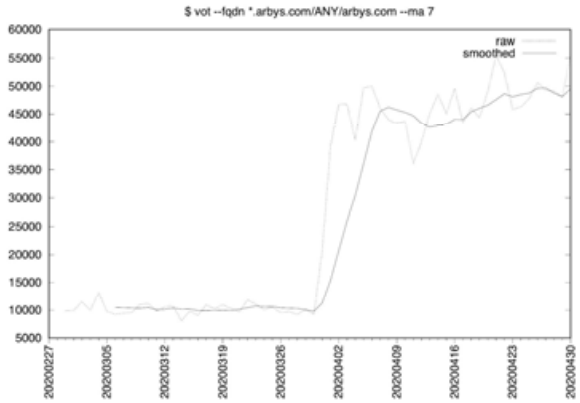

2. Burger King

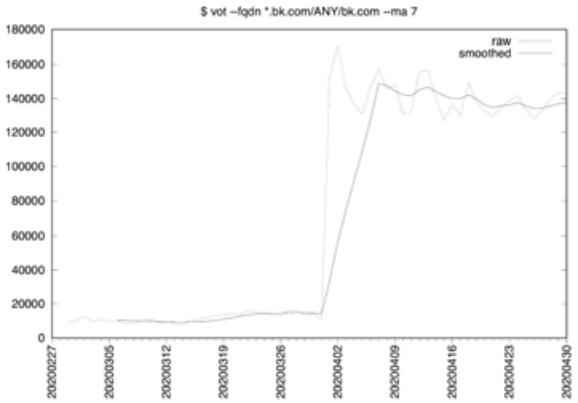

3. Chick-Fil-A

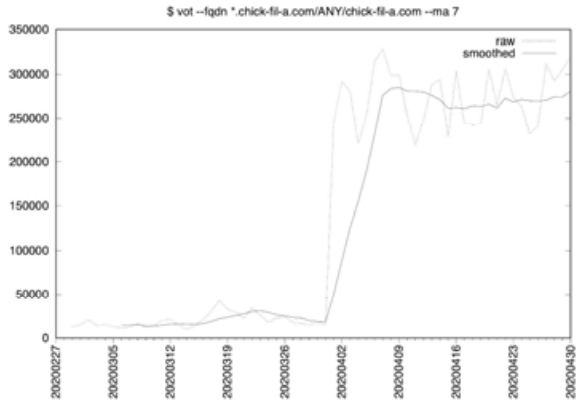

4. McDonald's

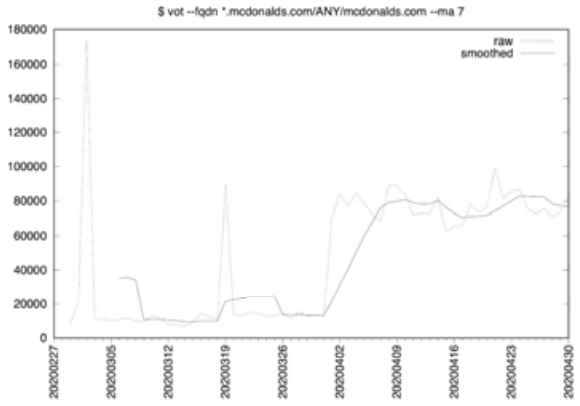

5. Starbucks [ATYPICAL SHAPE]

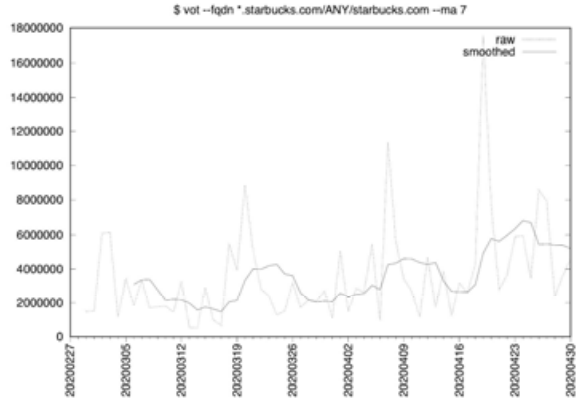

6. Subway

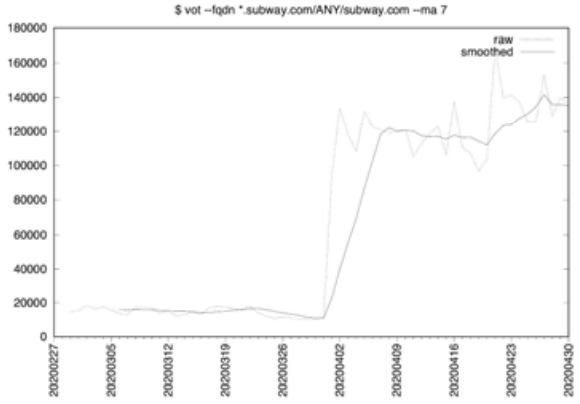

7. Taco Bell

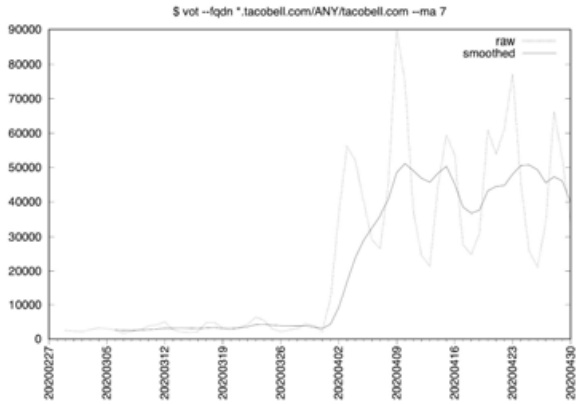

8. Wendy's [ATYPICAL SHAPE]

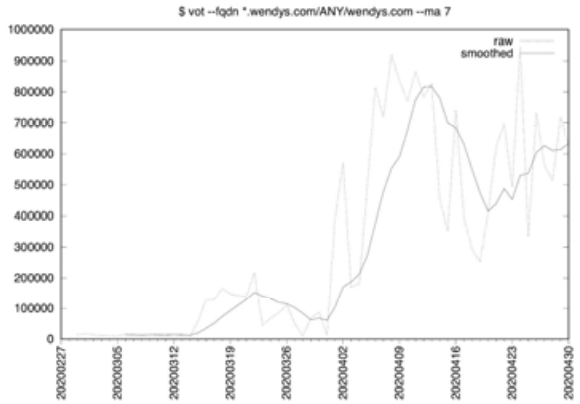

Supermarkets and Discount Club Stores

1. albertsons.com
2. bjs.com
3. costco.com
4. foodlion.com
5. heb.com
6. kroger.com
7. meijer.com
8. publix.com
9. safeway.com
10. samsclub.com
11. stopandshop.com
12. tesco.com

1. Albertsons

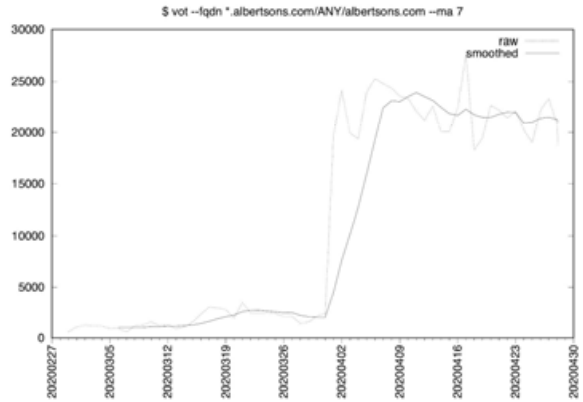

2. BJ's Wholesale Club

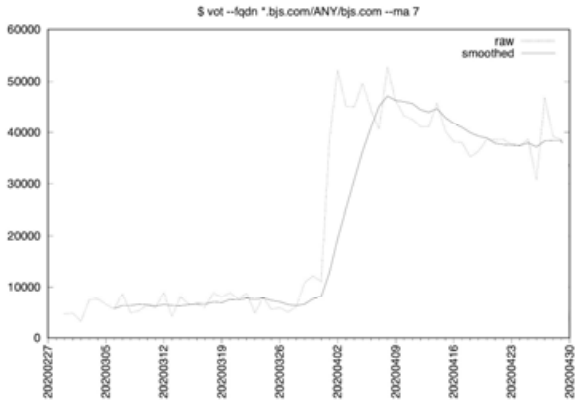

3. Costco

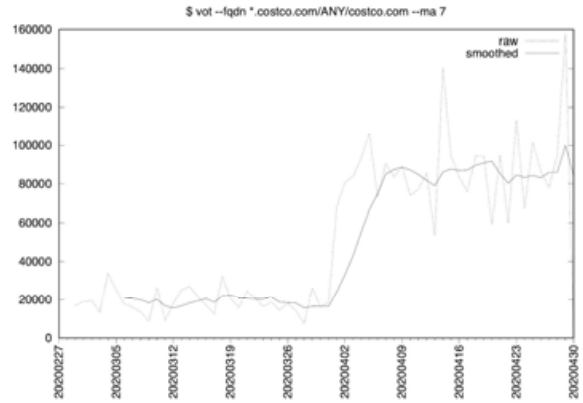

4. Food Lion

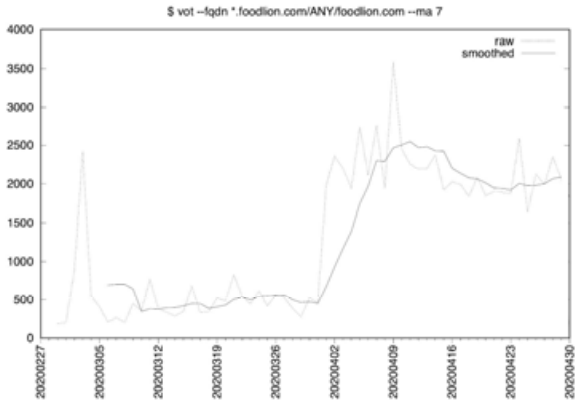

5. H-E-B

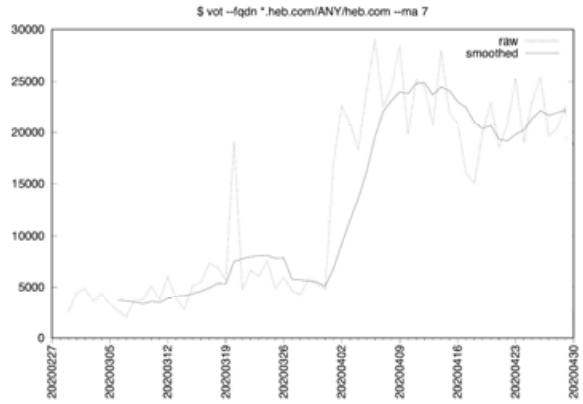

6. Kroger

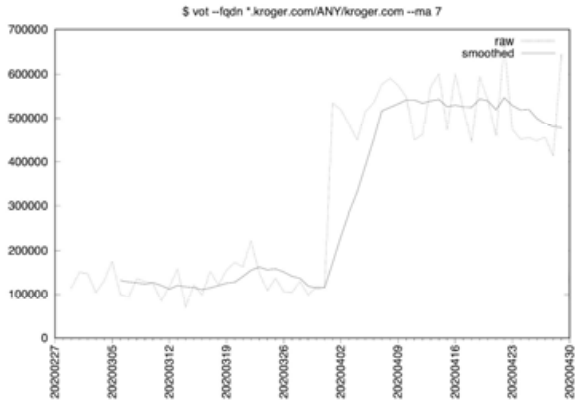

7. Meijer [ATYPICAL SHAPE]

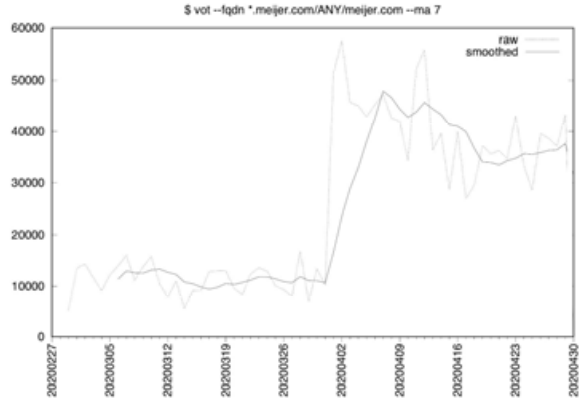

8. Publix [ATYPICAL SHAPE]

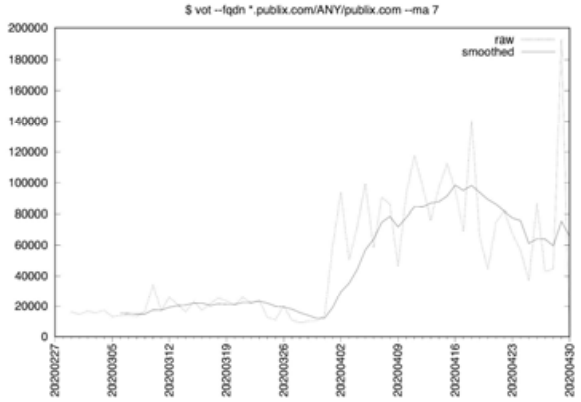

9. Safeway [ATYPICAL SHAPE]

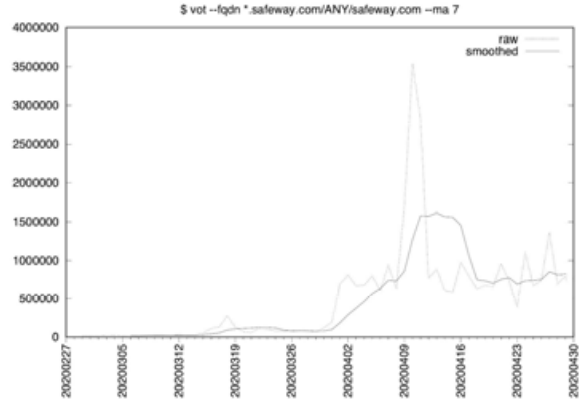

10. Sam's Club [ATYPICAL SHAPE]

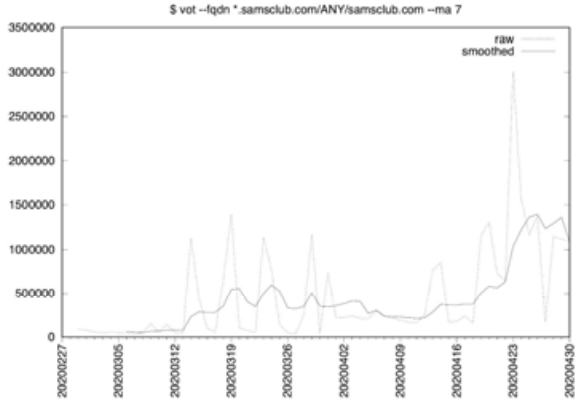

11. Stop and Shop [ATYPICAL SHAPE]

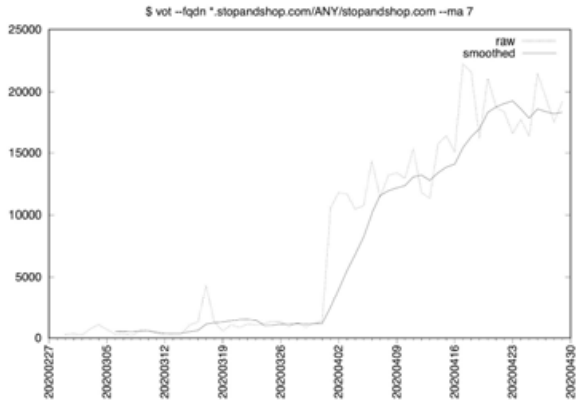

12. Tesco [ATYPICAL SHAPE]

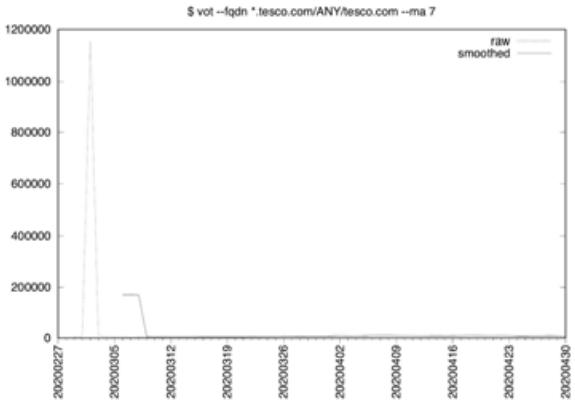

Streaming Video Sites

1. Hulu

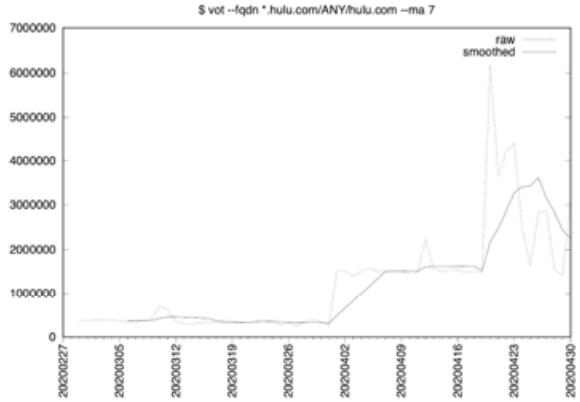

2. Netflix [ATYPICAL SHAPE]

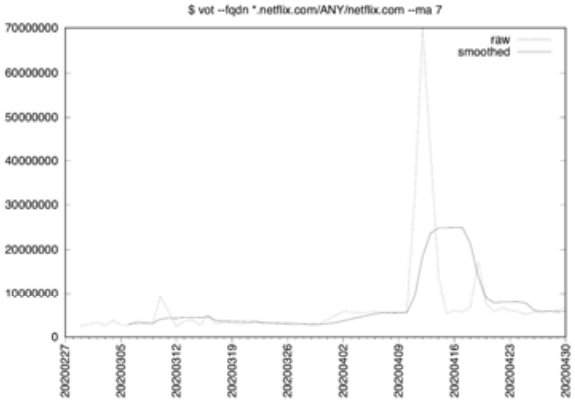

3. YouTube

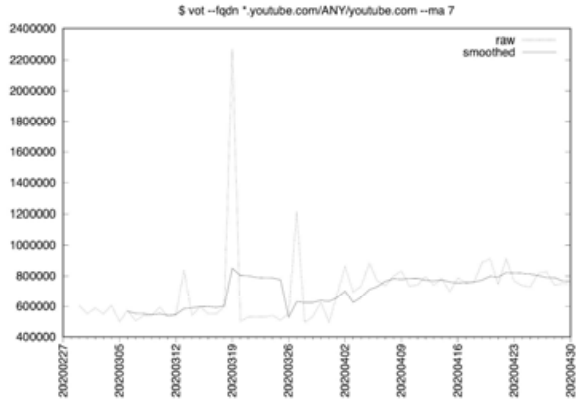

Higher Education Sites

NOTE: For the purposes of this section, the "typical" shape we've identified is different than the "typical" shape for most other sites in this report. For higher education, the "typical" shape is one that gradually rises to a "hill" then descends to a new value somewhere between the original level and the passed peak (e.g., the initial region may gradually ramp, and the new "plateau" tends not to be maintained).

Public Research Universities (Alphabetized By Domain Name)

1. arizona.edu
2. asu.edu
3. berkeley.edu
4. colorado.edu
5. fsu.edu
6. indiana.edu
7. msu.edu
8. ncsu.edu
9. olemiss.edu
10. oregonstate.edu
11. osu.edu
12. psu.edu
13. tamu.edu
14. ua.edu
15. uaf.edu
16. ucla.edu
17. ucsb.edu
18. ucsd.edu
19. ufl.edu
20. uga.edu
21. uiowa.edu
22. uiuc.edu
23. umd.edu
24. umich.edu
25. umn.edu
26. unc.edu
27. uoregon.edu
28. utexas.edu
29. vt.edu
30. washington.edu
31. wisc.edu

1. University of Arizona

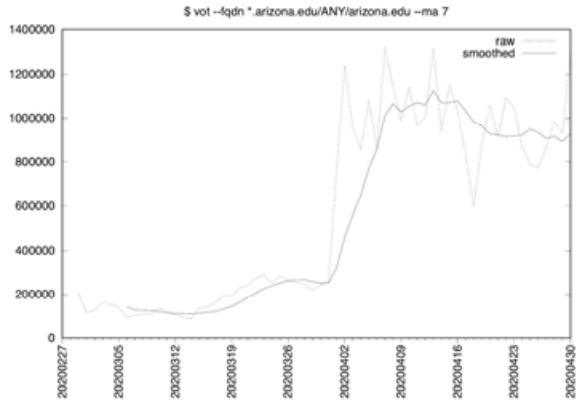

2. Arizona State University [ATYPICAL SHAPE]

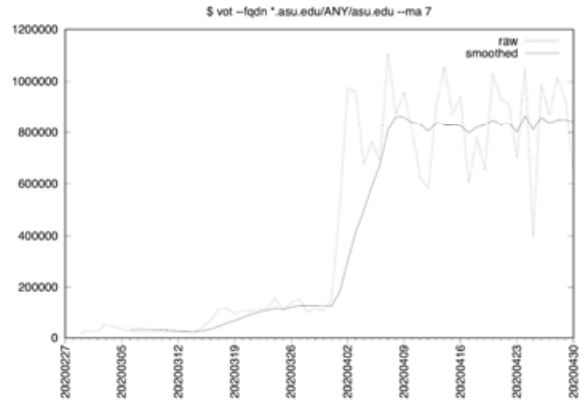

3. Berkeley [ATYPICAL SHAPE]

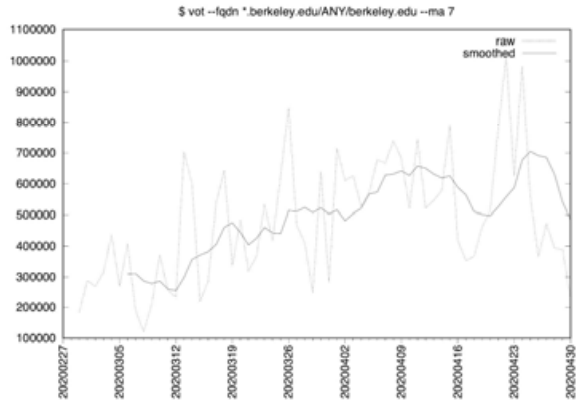

4. University of Colorado

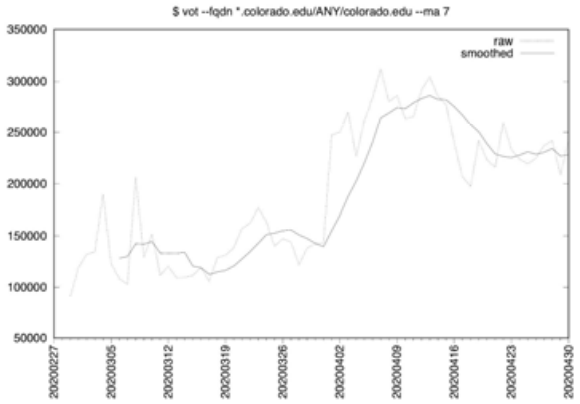

5. Florida State University

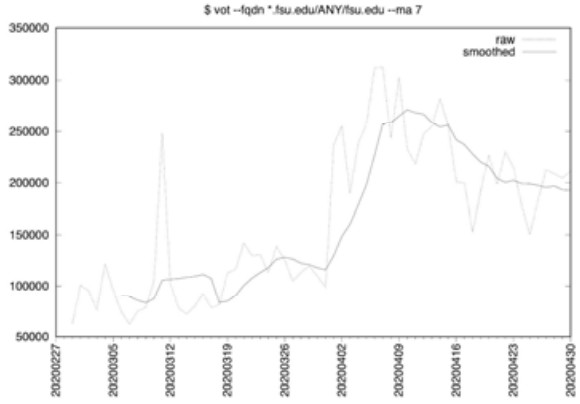

6. Indiana University

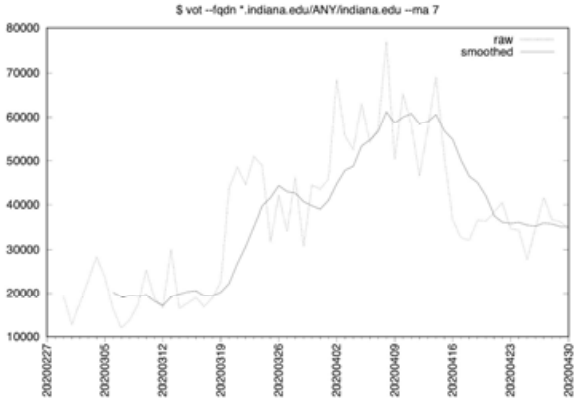

7. Michigan State University

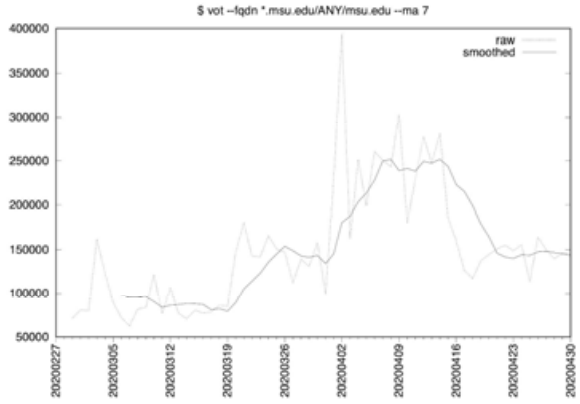

8. North Carolina State University

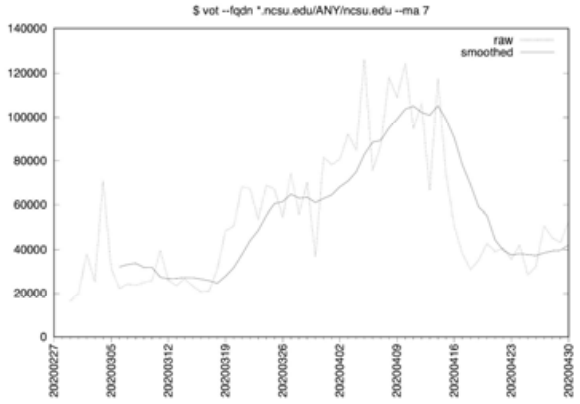

9. University of Mississippi

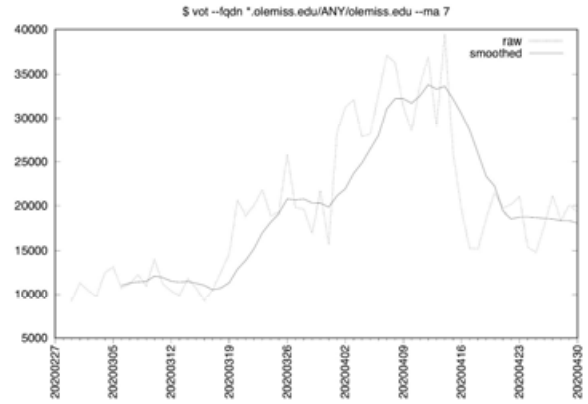

10. Oregon State University

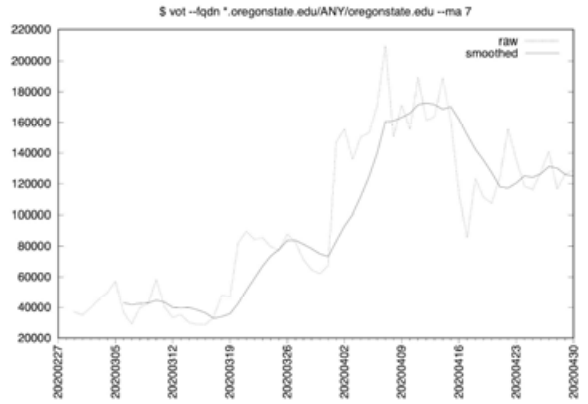

11. Ohio State University

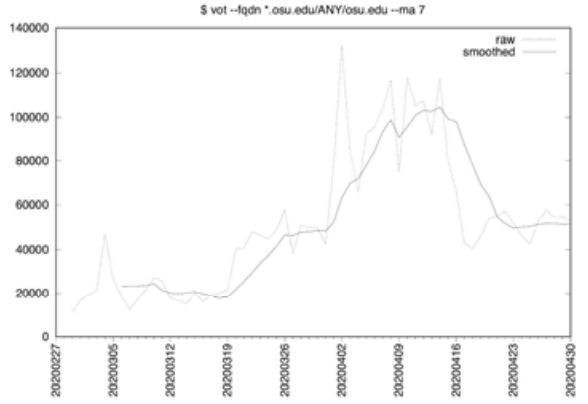

12. Penn State University

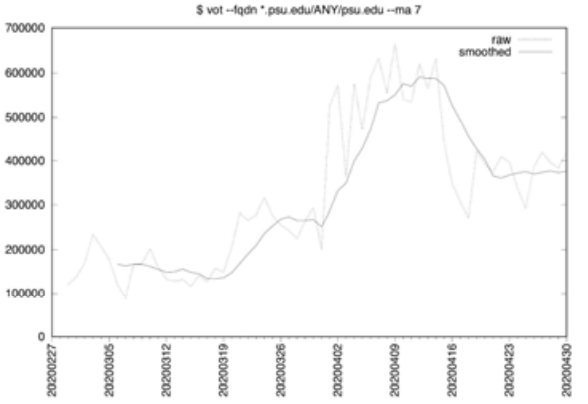

13. Texas A&M University

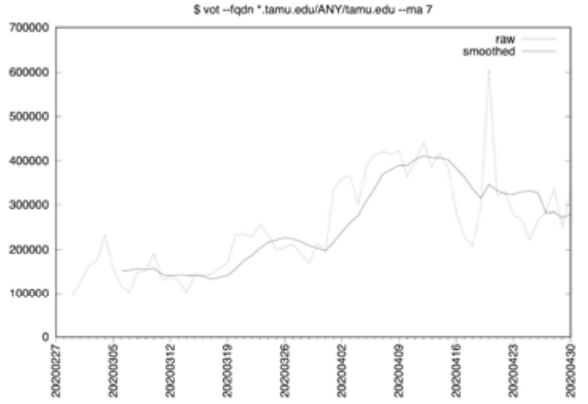

14. University of Alabama

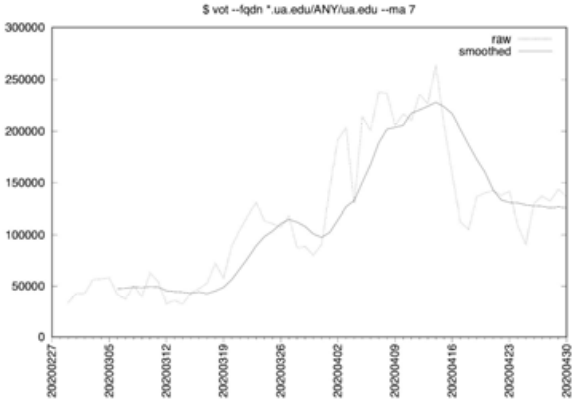

15. University of Alaska Fairbanks

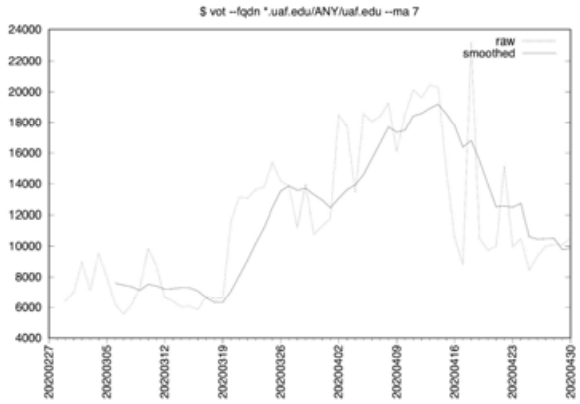

16. University of California Los Angeles [ATYPICAL SHAPE]

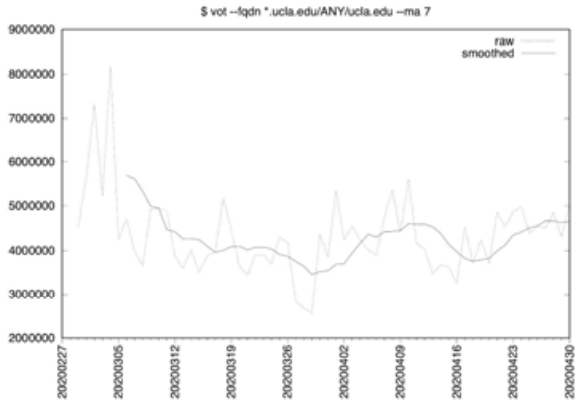

17. University of California Santa Barbara

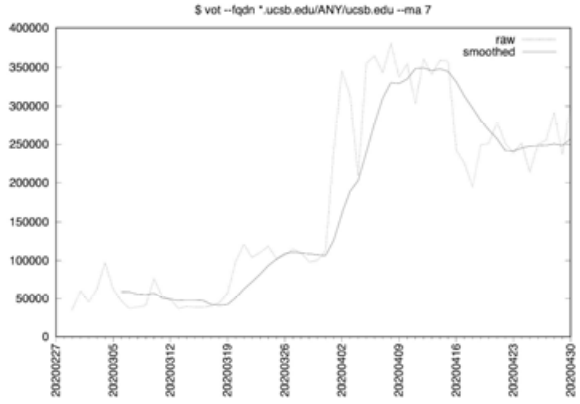

18. University of California San Diego

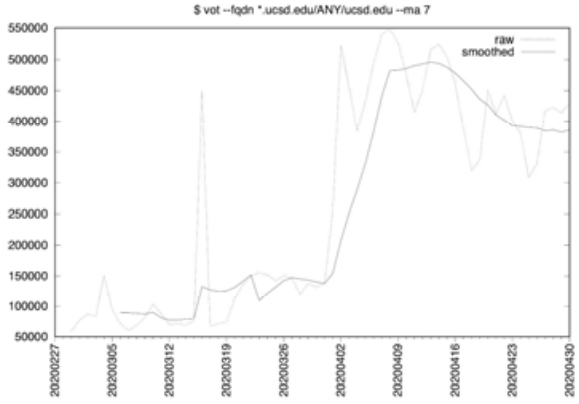

19. University of Florida

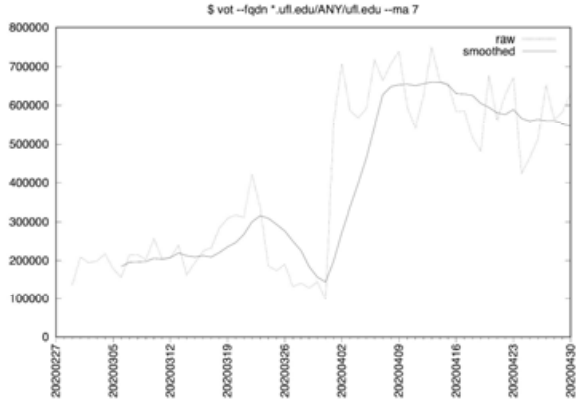

20. University of Georgia

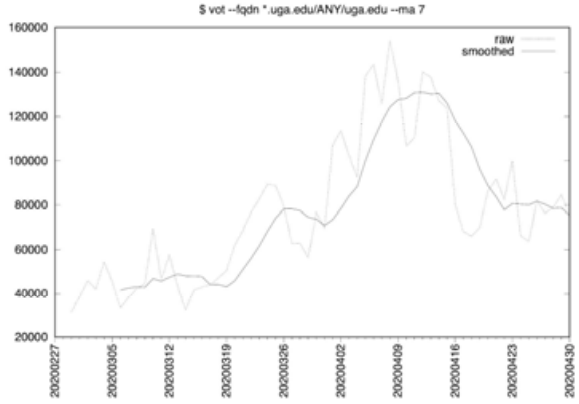

21. University of Iowa

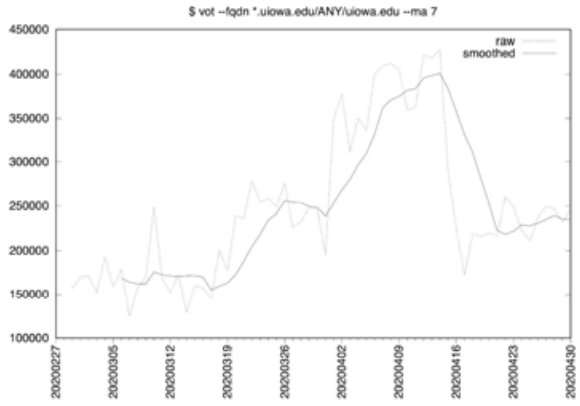

22. University of Illinois Urbana Champaign

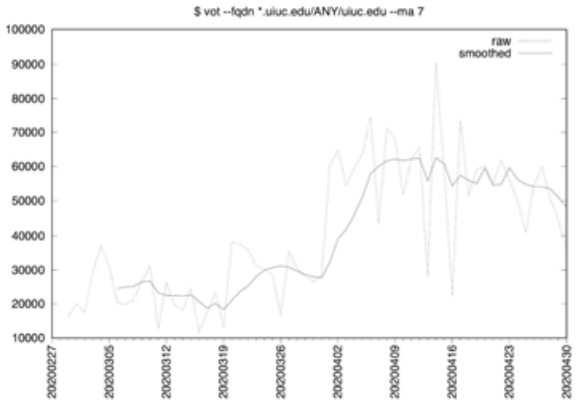

23. University of Maryland

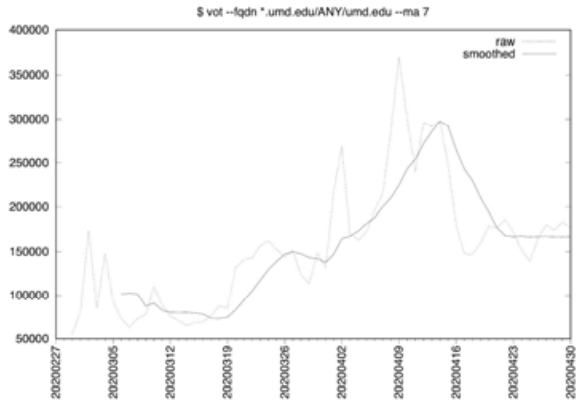

24. University of Michigan

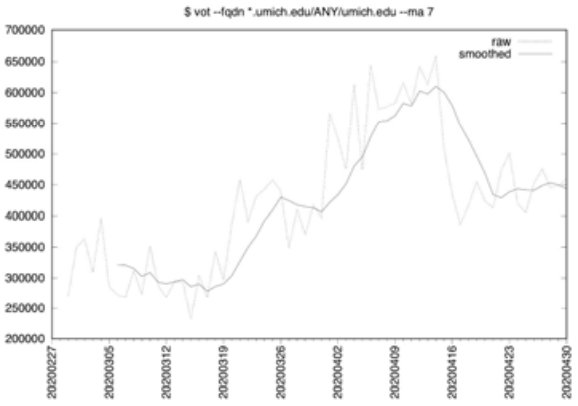

25. University of Minnesota

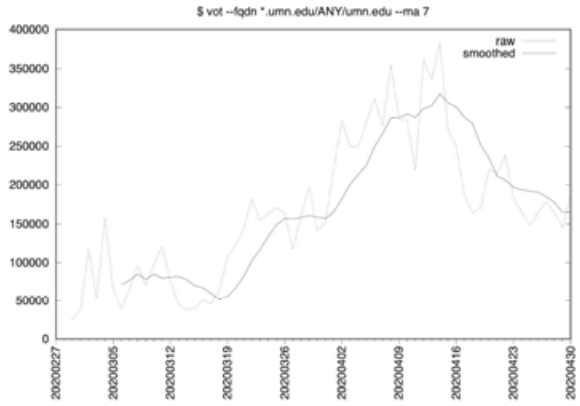

26. University of North Carolina

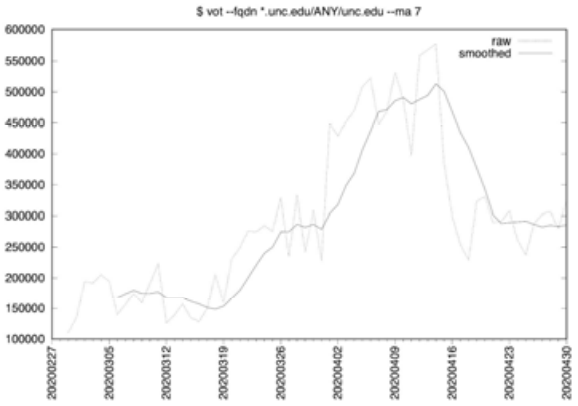

27. University of Oregon

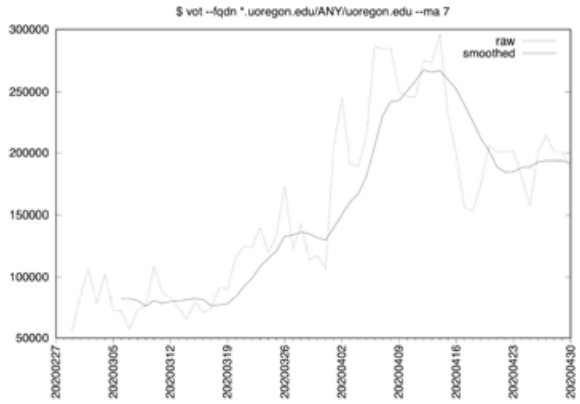

28. University of Texas

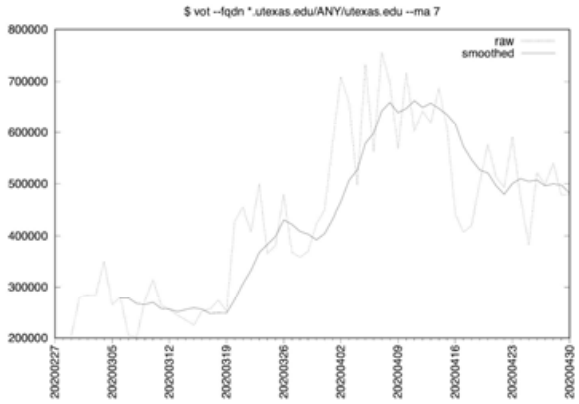

29. Virginia Tech

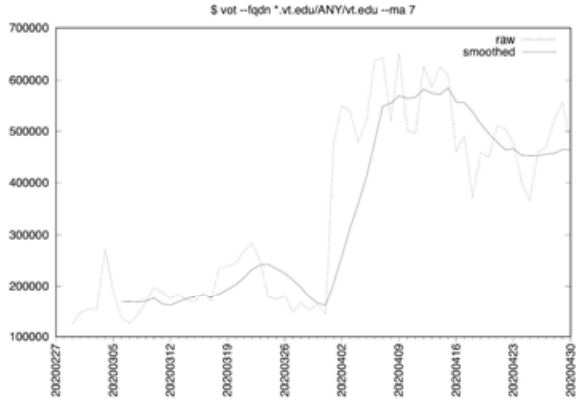

30. University of Washington [ATYPICAL SHAPE]

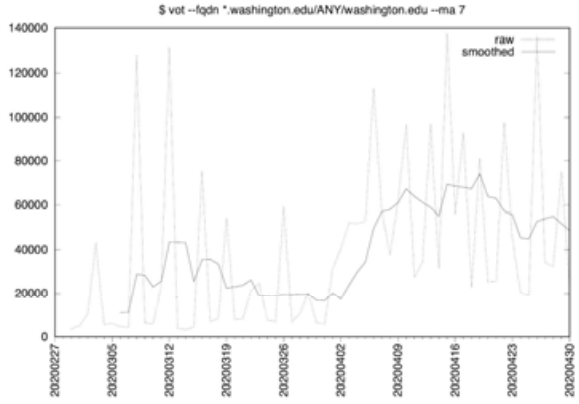

31. University of Wisconsin

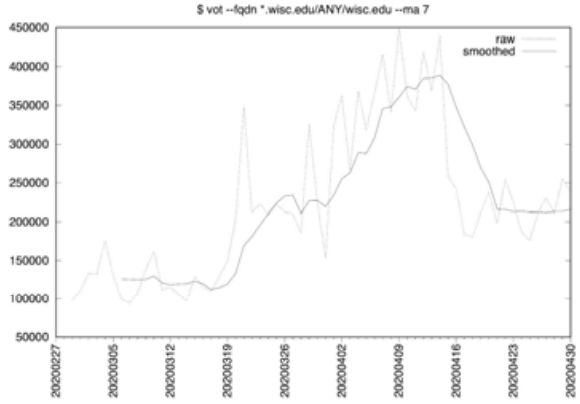

Private Research Universities (Alphabetized By Domain Name)

1. bu.edu
2. byu.edu
3. caltech.edu
4. clemson.edu
5. cmu.edu
6. duke.edu
7. emory.edu
8. gatech.edu
9. georgetown.edu
10. gwu.edu
11. jhu.edu
12. mit.edu
13. nd.edu
14. northwestern.edu
15. nyu.edu
16. pitt.edu
17. rice.edu
18. rpi.edu
19. stanford.edu
20. tufts.edu
21. uchicago.edu
22. usc.edu
23. vanderbilt.edu
24. virginia.edu
25. wfu.edu
26. wustl.edu
27. wm.edu

1. Boston University [ATYPICAL SHAPE]

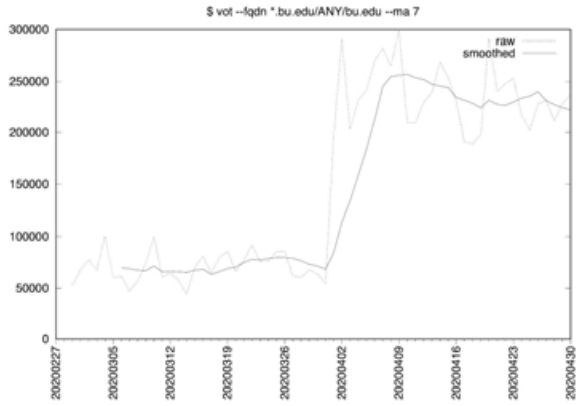

2. Brigham Young University

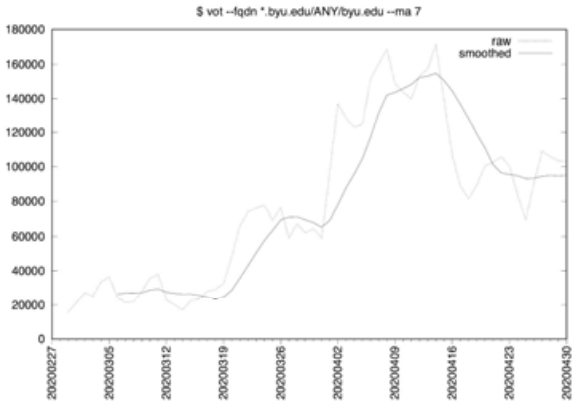

3. California Institute of Technology

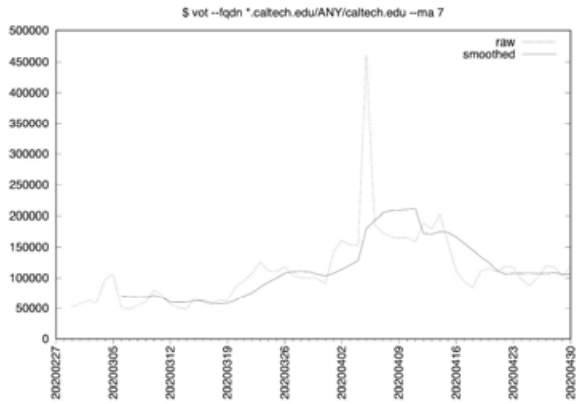

4. Clemson University [ATYPICAL SHAPE]

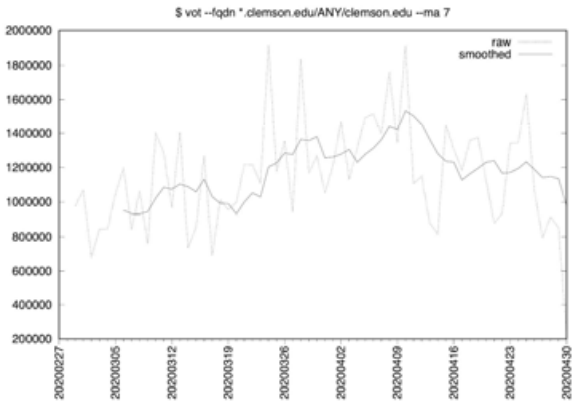

5. Carnegie Mellon University

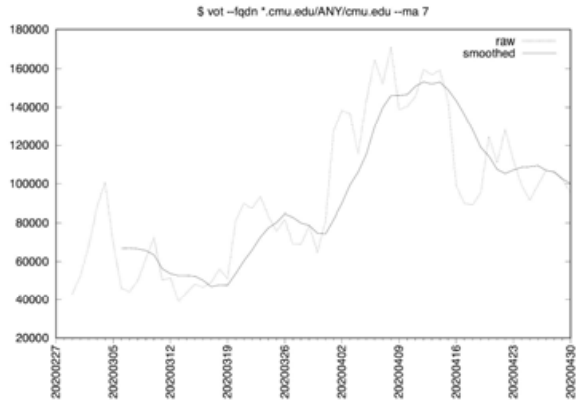

6. Duke

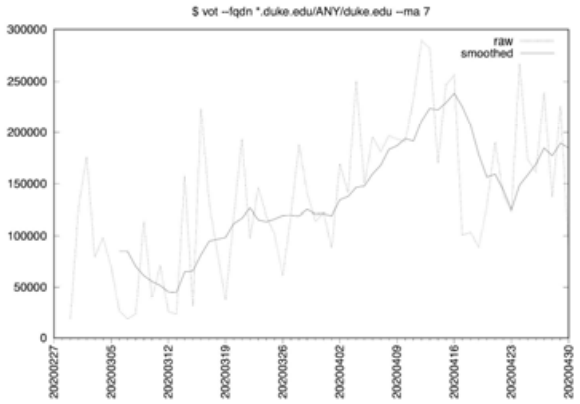

7. Emory University

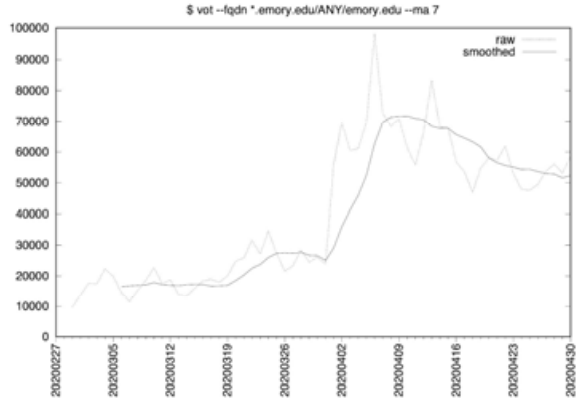

8. Georgia Tech University

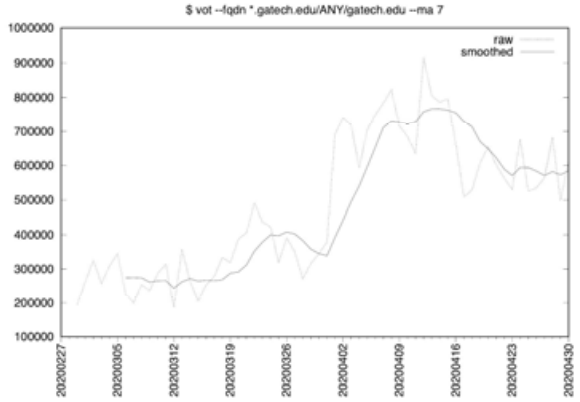

9. Georgetown University [ATYPICAL SHAPE]

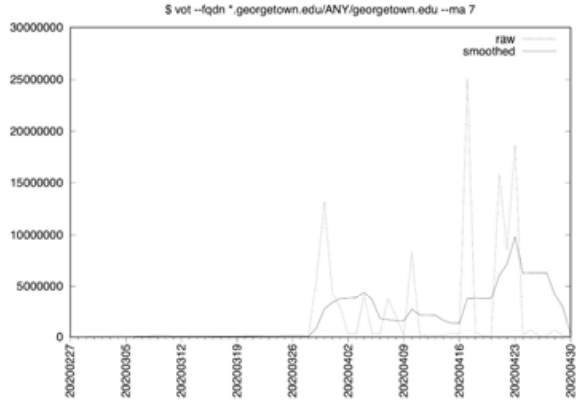

10. George Washington University

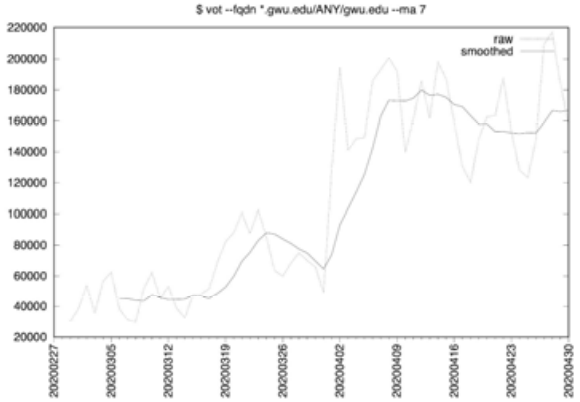

11. Johns Hopkins University

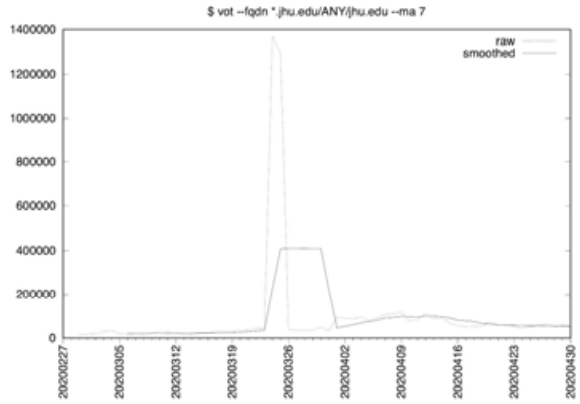

12. MIT

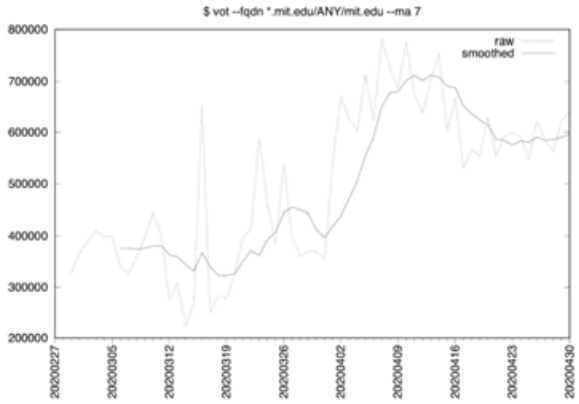

13. Notre Dame University

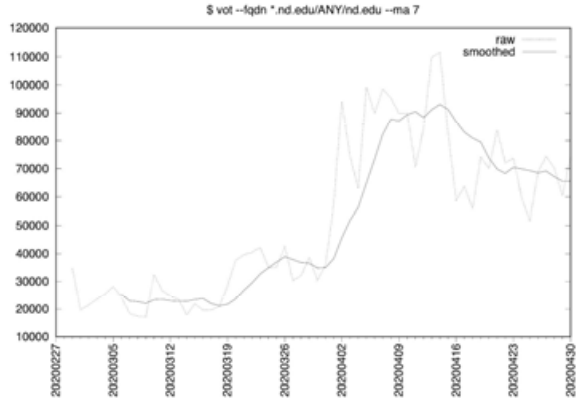

14. Northwestern University

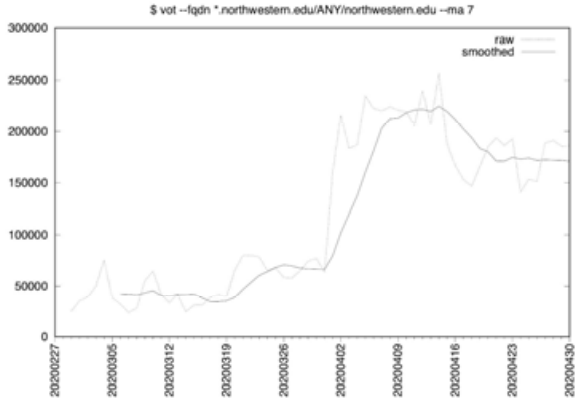

15. NYU [ATYPICAL SHAPE]

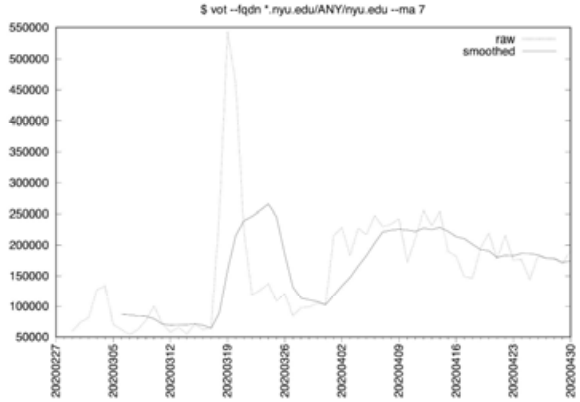

16. University of Pittsburgh

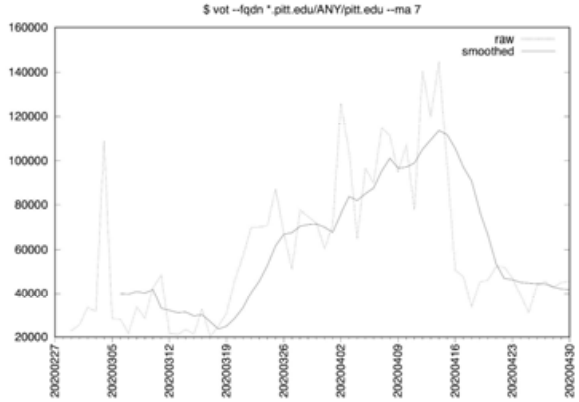

17. Rice University

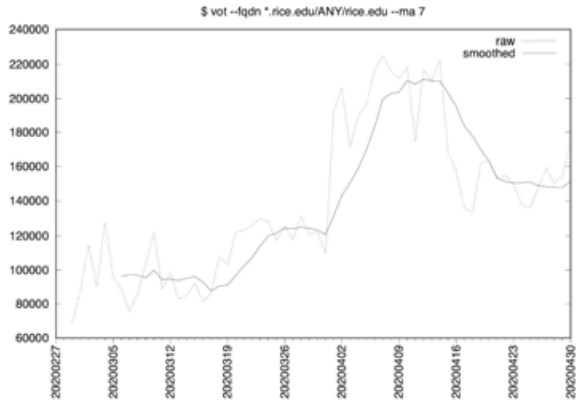

18. Rensellaer Polytechnic Institute (RPI)

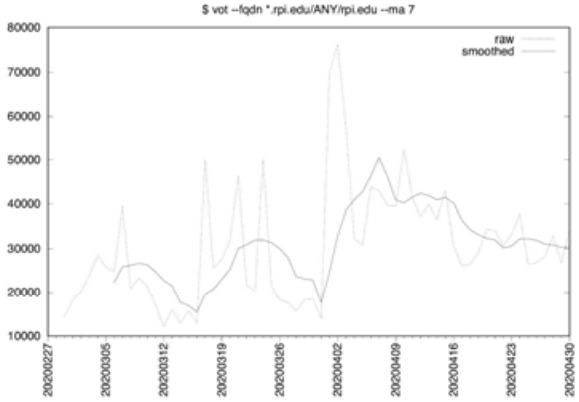

19. Stanford University

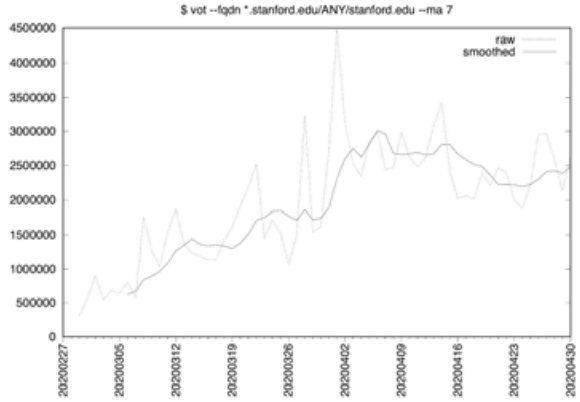

20. Tufts

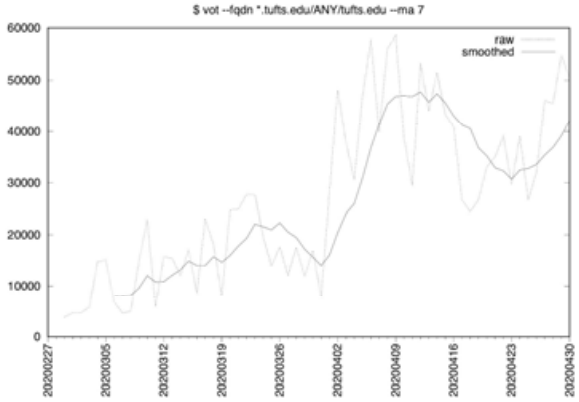

21. University of Chicago

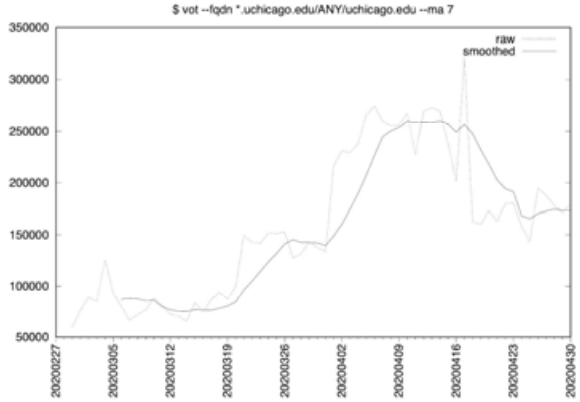

22. University of Southern California

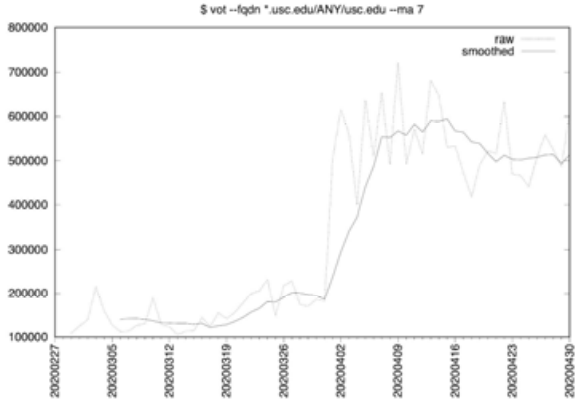

23. Vanderbilt

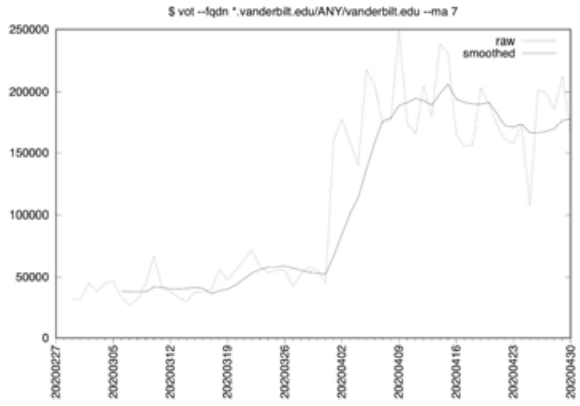

24. University of Virginia

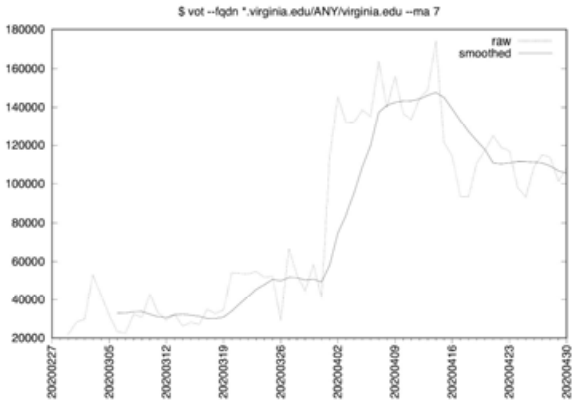

25. Wake Forest University

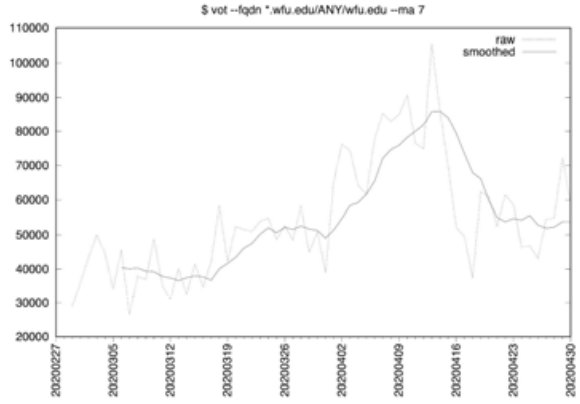

26. Washington University of St Louis

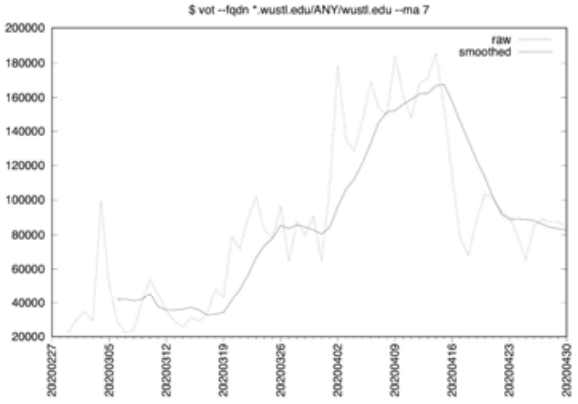

27. William and Mary

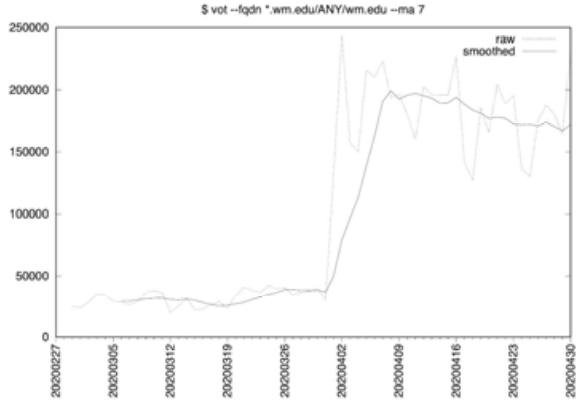

The Ivy League (Alphabetized by Domain Name)

1. brown.edu
2. columbia.edu
3. cornell.edu
4. dartmouth.edu
5. harvard
6. princeton.edu
7. upenn.edu
8. yale.edu

1. Brown

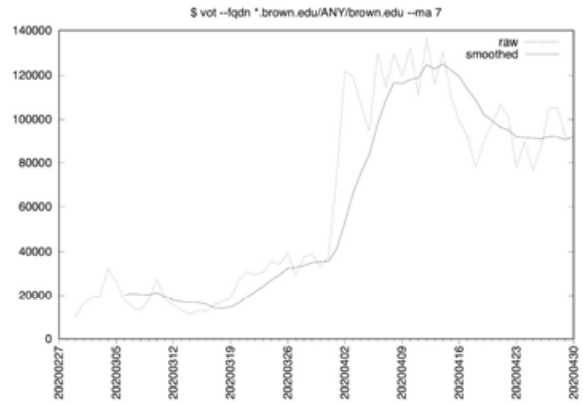

2. Columbia University

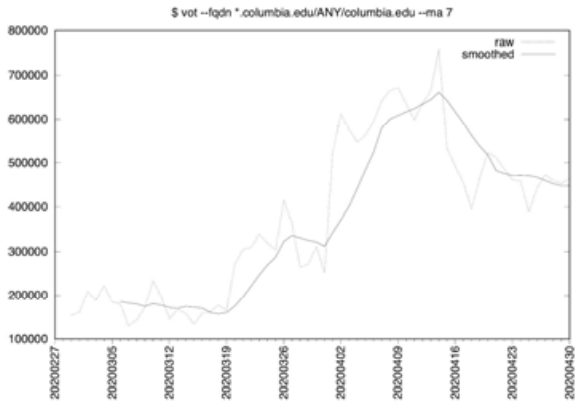

3. Cornell University

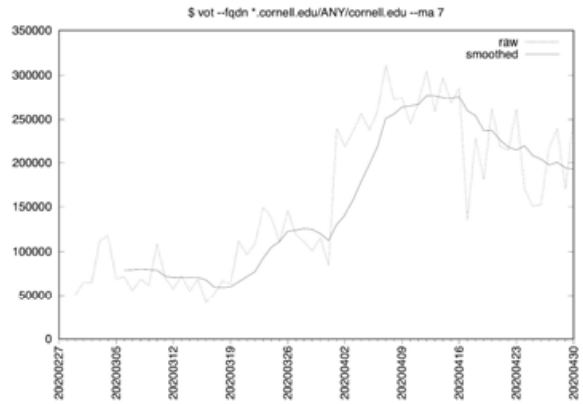

4. Dartmouth

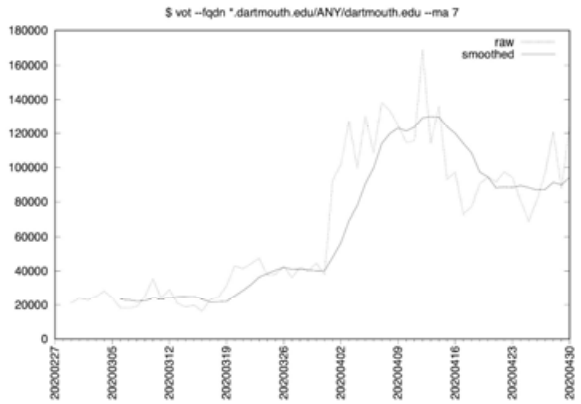

5. Harvard

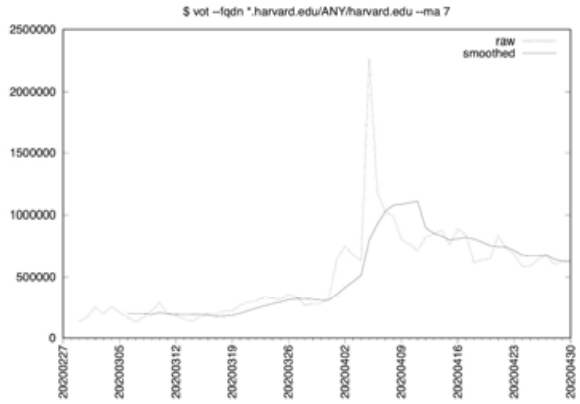

6. Princeton

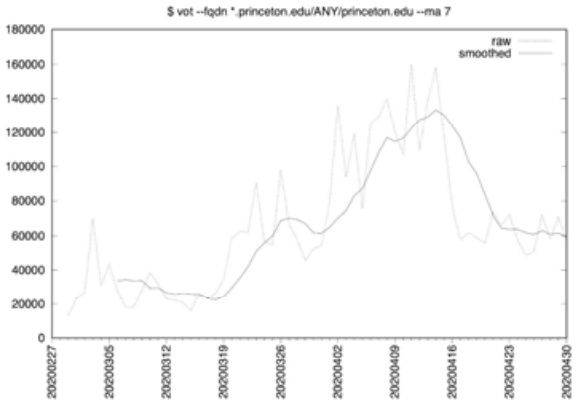

7. University of Pennsylvania

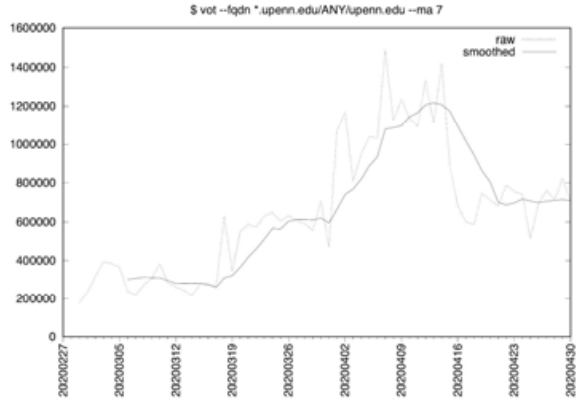

8. Yale

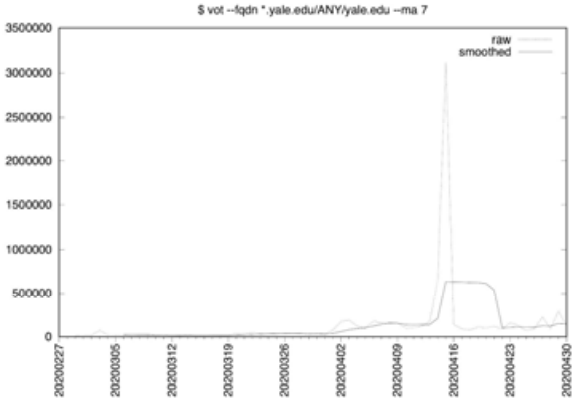

Liberal Arts Colleges (Alphabetized by Domain Name)

1. amherst.edu
2. bowdoin.edu
3. carleton.edu
4. colby.edu
5. grinnell.edu
6. hmc.edu
7. pomona.edu
8. reed.edu
9. wesleyan.edu
10. whitman.edu
11. williams.edu

1. Amherst

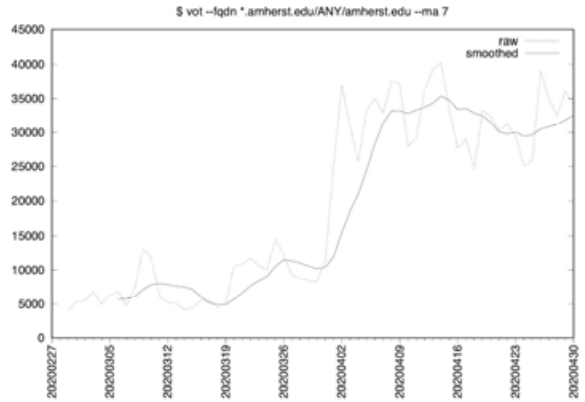

2. Bowdoin

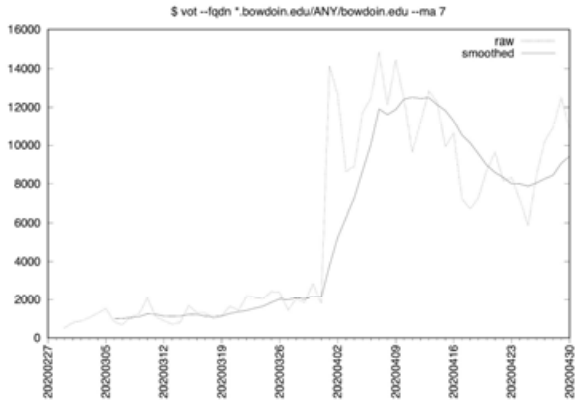

3. Carleton College

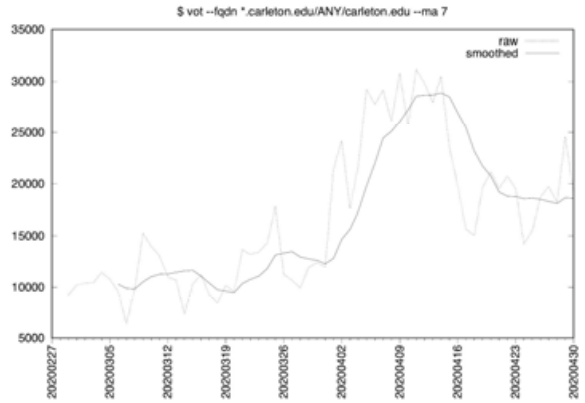

4. Colby College

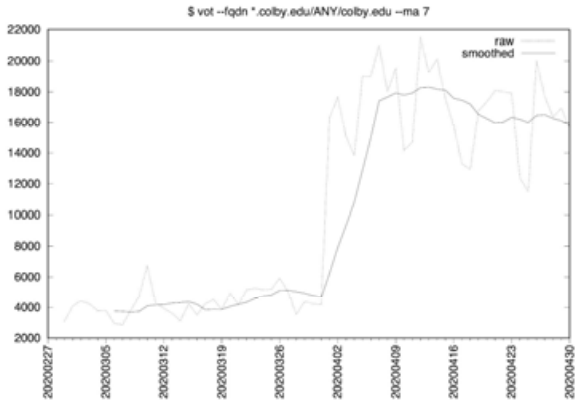

5. Grinnell College [ATYPICAL SHAPE]

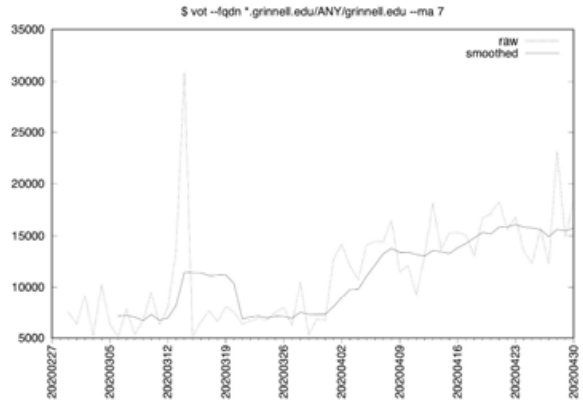

6. Harvey Mudd College

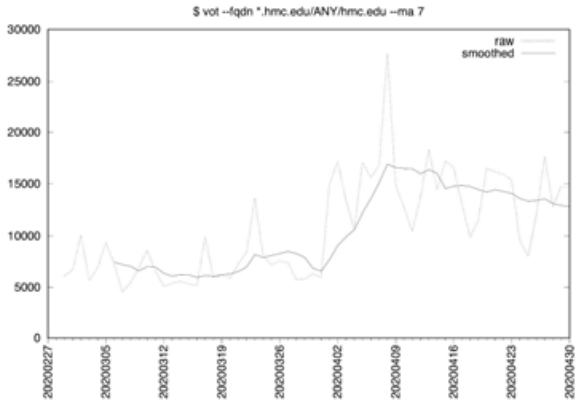

7. Pomona College

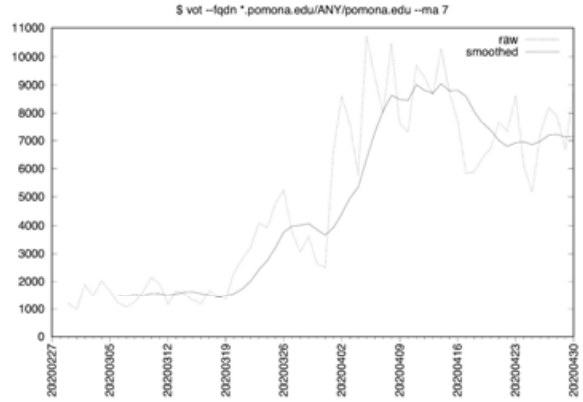

8. Reed College

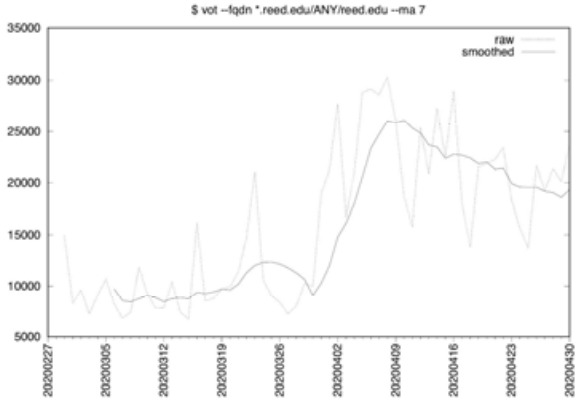

9. Wesleyan University

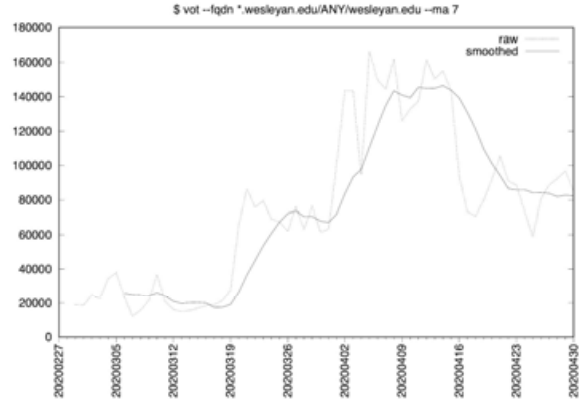

10. Whitman College

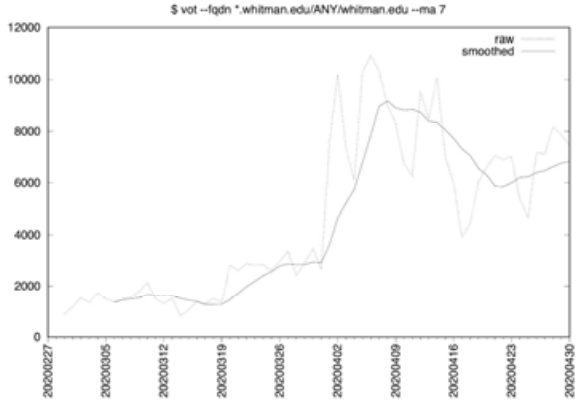

11. Williams College

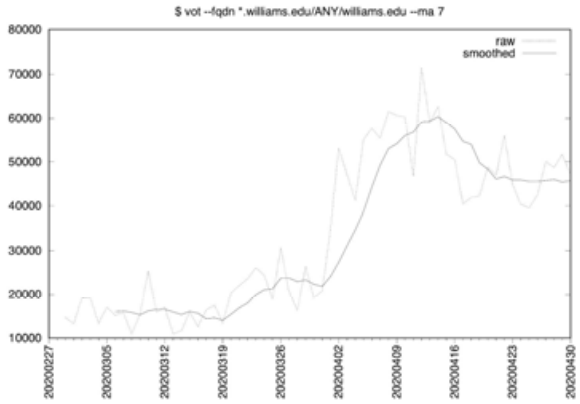

International Universities (Alphabetized by Domain Name)

1. anu.edu.au
2. cam.ac.uk
3. ed.ac.uk
4. epfl.ch
5. ethz.ch
6. imperial.ac.uk
7. kcl.ac.uk
8. lse.ac.uk
9. mcgill.ca
10. nus.edu.sg
11. ox.ac.uk
12. pku.edu.cn
13. tsinghua.edu.cn
14. tum.de
15. u-tokyo.ac.jp
16. ubc.ca
17. ucl.ac.uk
18. uni-heidelberg.de
19. uni-muenchen.de
20. unimelb.edu.au
21. utoronto.ca

1. Australian National University, Australia

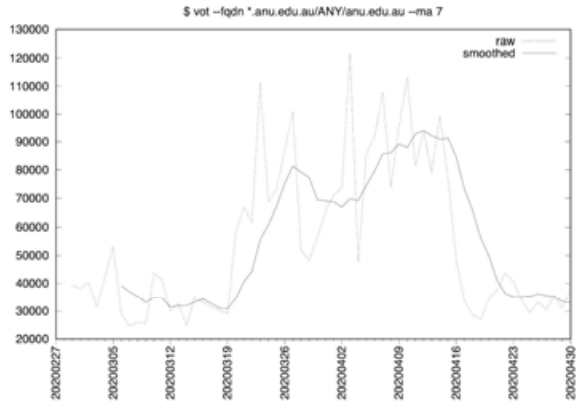

2. Cambridge University, England [ATYPICAL SHAPE]

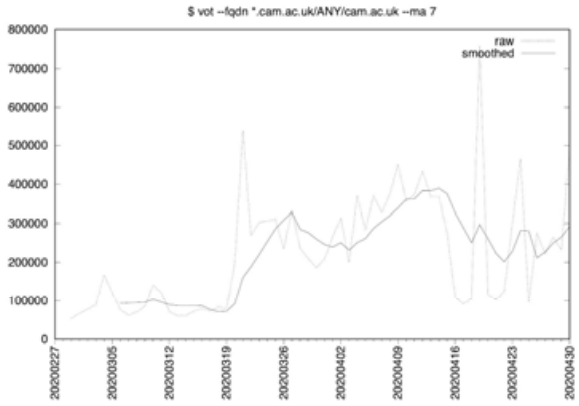

3. University of Edinburgh, Scotland [ATYPICAL SHAPE]

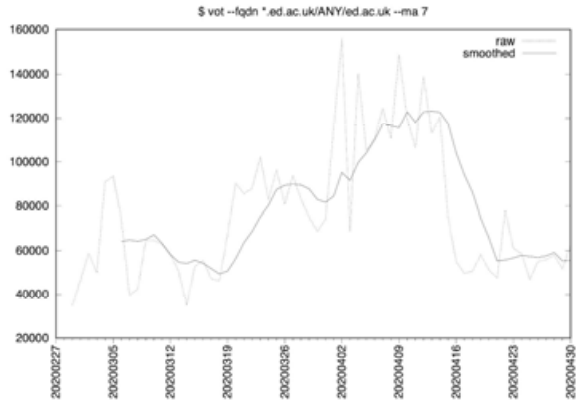

4. EPFL, Switzerland [ATYPICAL SHAPE]

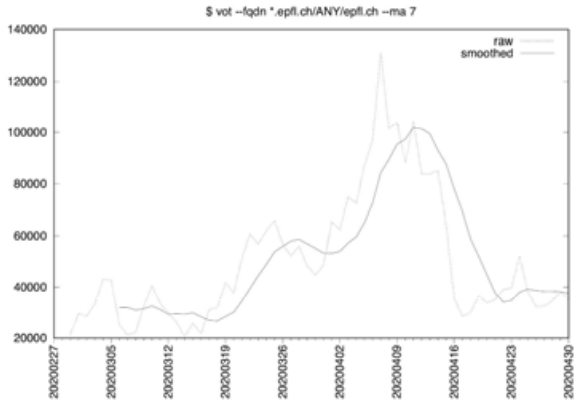

5. ETHZ, Switzerland

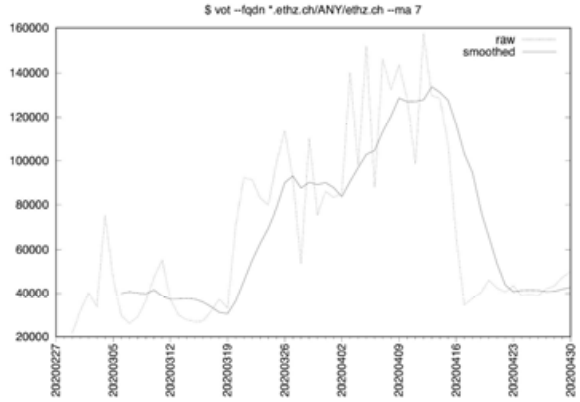

6. Imperial College, England

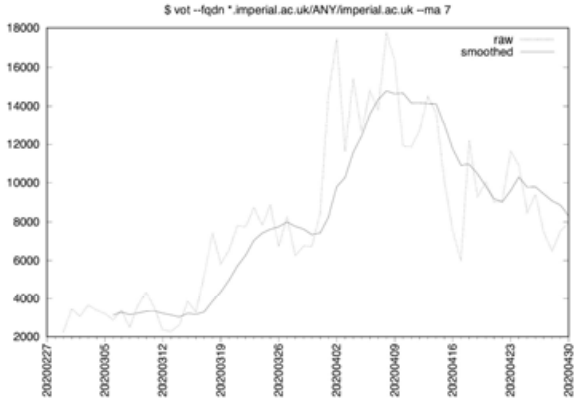

7. King's College London, England

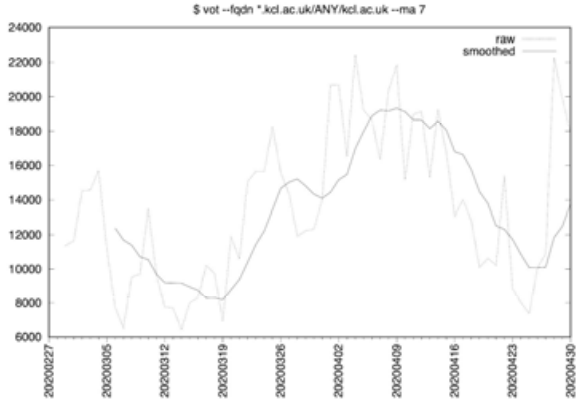

8. London School of Economics, England

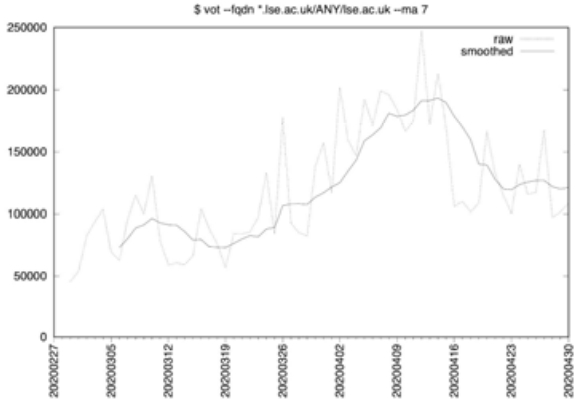

9. McGill University, Canada

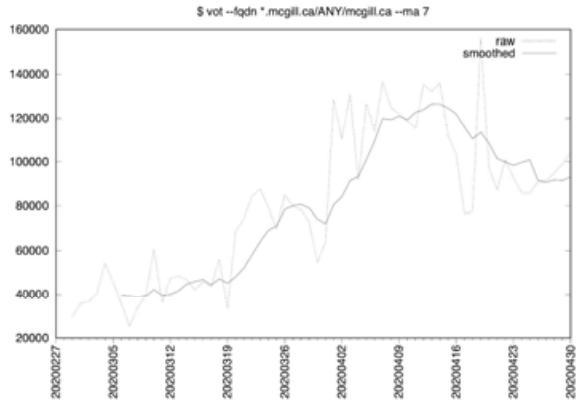

10. National University of Singapore [ATYPICAL SHAPE]

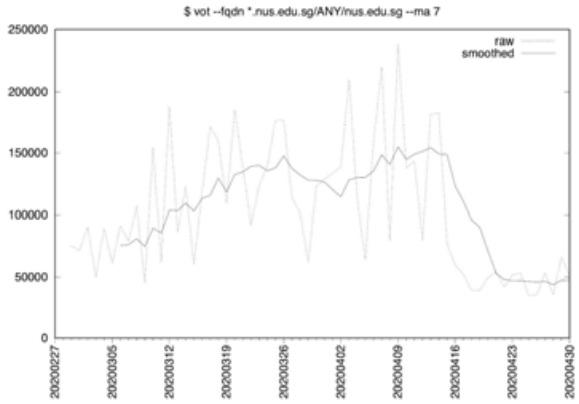

11. University of Oxford, England

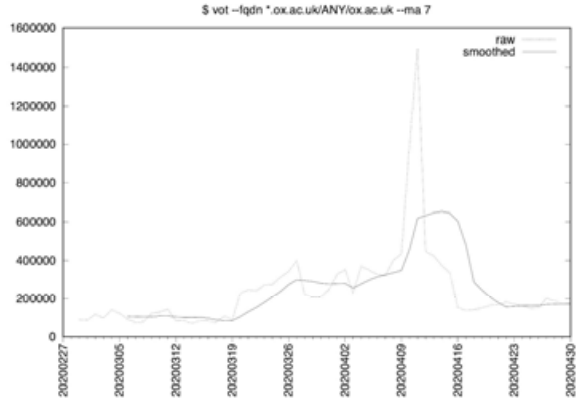

12. Peking University, China [ATYPICAL SHAPE]

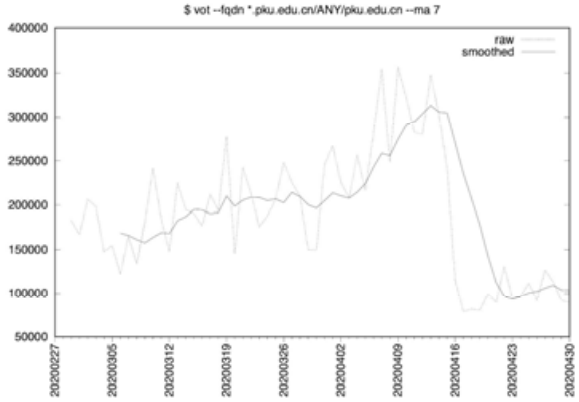

13. Tsinghua University, China [ATYPICAL SHAPE]

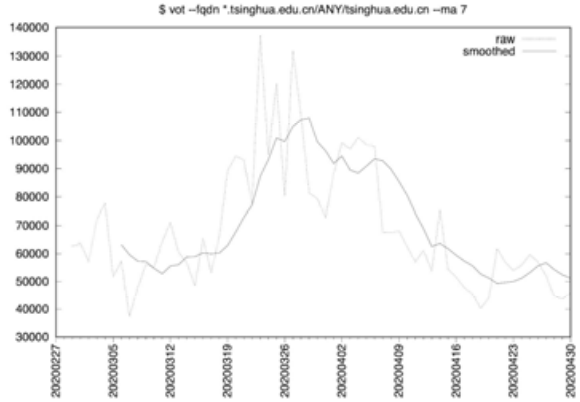

14. Technical University of Munich, Germany

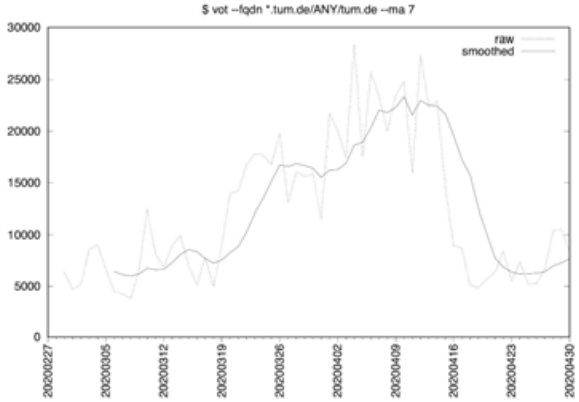

15. University of Tokyo, Japan

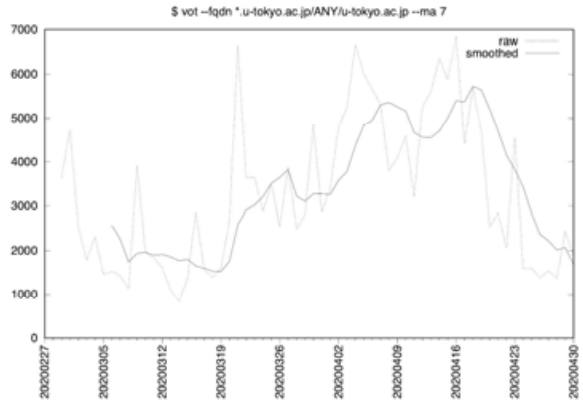

16. University of British Columbia, Canada

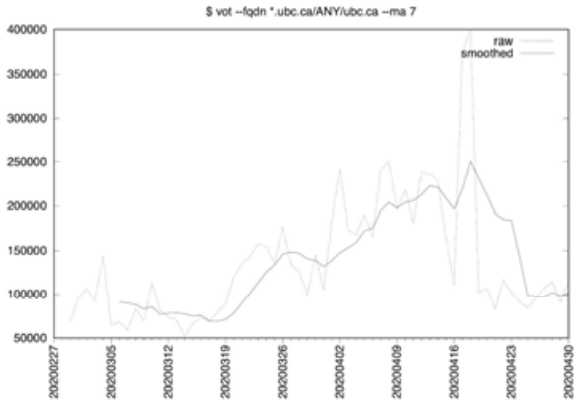

17. University College London, England

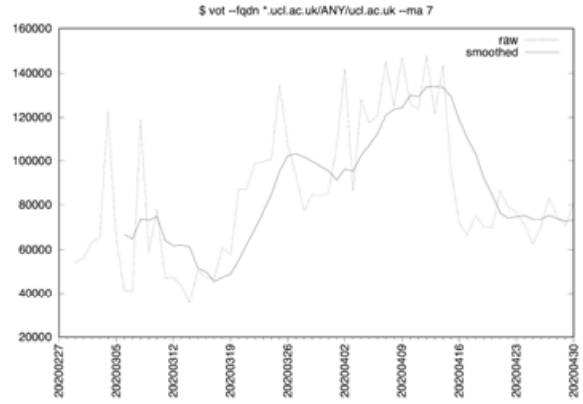

18. University of Heidelberg, Germany

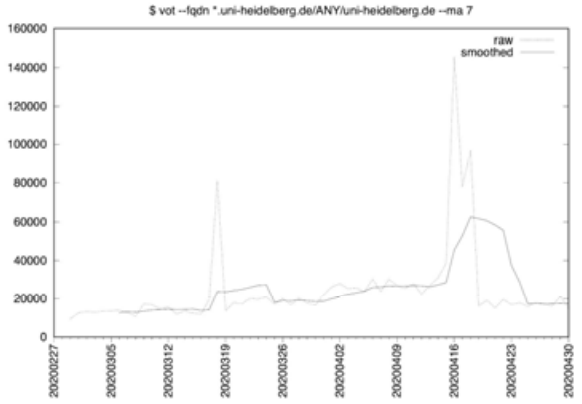

19. University of Munich, Germany

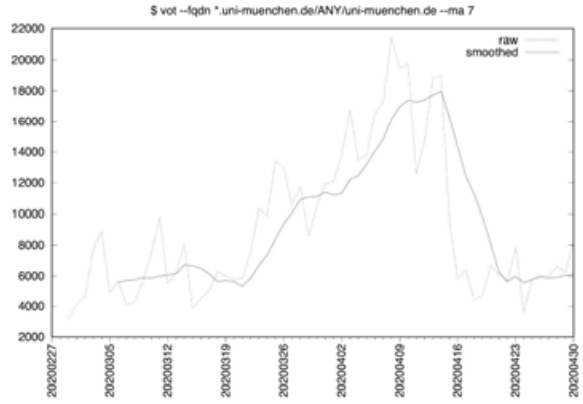

20. University of Melbourne, Australia

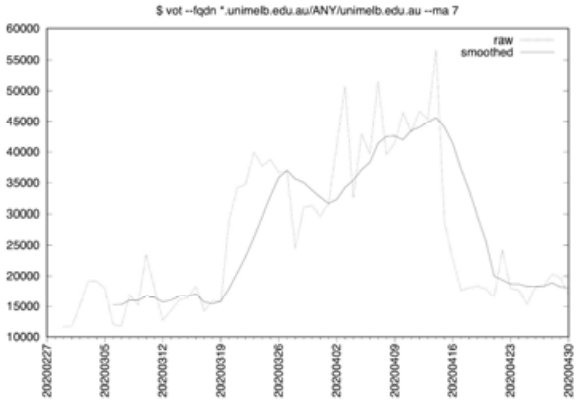

21. University of Toronto, Canada

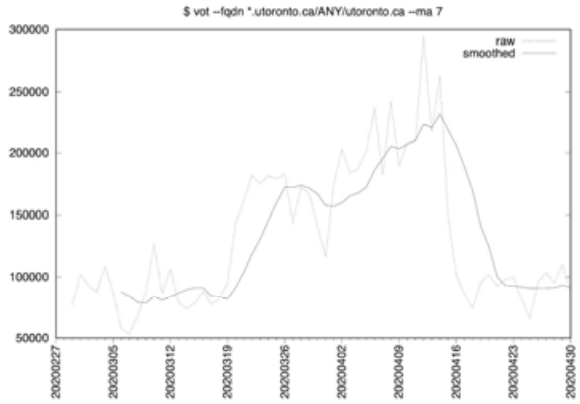

Section III. Conclusion and Potential Future Work

Conclusion: You've how seen how DNS cache miss traffic volumes have risen during the COVID-19 pandemic of 2020 across selected industries: generally speaking, we see a "step up" pattern typically reflecting a 4x-to-7x increase in DNS cache miss traffic levels.

This change took place, often abruptly, between mid-to-late-March and early-to-mid-April. While most of the studied sites exhibited this characteristic traffic pattern, there was variation among the studied sites in terms of magnitude and timing, and higher education sites tends to exhibit an increase, but that increase would then subsequently drop, producing a hill rather than a plateau.

We provide individual graphs for each of the 316 sites we studied so you can see what we observed yourself.

While examining that traffic, it was impossible to miss some sites where "spike" patterns dominated normal traffic levels. These spikes appear to be denial of service attacks consisting of two main types:

- One type that appears to be purely associated with abusive DNS SOA ("Start of Authority") query levels, and
- A second type that melds abusive DNS SOA query levels with abusive DNS TXT queries for wildcarded SPF redirect records.

Farsight recommends that name server vendors ship their products with Response Rate Limiting (RRL) enabled by default. Farsight also recommends that all authoritative name server operators confirm that their current configurations have RRL enabled.

Potential Future Work: While this report represents a substantial body of work in its own right, there are obvious opportunities to extend it, including:

Immediately Possible Work:

- *More/Different Sites:* We could expand the set of sites we consider, particularly when it comes to sites outside the United States. Rather than taking an industry-by-industry approach, we might want to consider just looking at the top <N> sites. We might also want to study the impact on health-related industries such as hospital systems, medical supply providers, pharmaceutical companies, etc.
- *Longer Study Period:* It is unlikely that COVID-19 will disappear any time soon (that will likely require development and widespread administration of a yet-to-be-available vaccine). Will currently elevated traffic levels continue, or will they gradually recede to normal levels?
- *Finer-Grained Analysis:* The current study pooled all FQDNs under each selected second level domain -- should we perhaps do a FQDN-by-FQDN study instead? If so, this might dictate an interactive atlas rather than a static report due to the sheer volume of graphs involved.
- *Focus on Anomalous Traffic Rather Than COVID-19-Related Impacts?* Should we look more closely for other evidence of malicious/anomalous traffic rather than focus on macroscopic volumetric changes? (We expected to see, and saw, an increase in traffic during the COVID-19 event; we did NOT expect to see the DDoS attack traffic we stumbled across)

Opportunities for Collaboration:

- *Correlation With Other Measurements:* There are numerous limitations associated with looking at cache miss traffic volumes. We'd be highly interested in seeing how our measurements relate to other traffic measurements, such as Netflow traffic measurements, or web site traffic volumes, on a site-by-site or FQDN-by-FQDN basis.

Future Work Requiring Changes in Farsight Data Collection:

- *TTL-Enhanced Analysis:* During the study it became clear that it would have been handy to have had TTL values available for each tuple (conceptually imagine the current RRname, RRtype, Bailiwick, Rdata tuple extended to also include a TTL value). If TTLs were tracked on a per-name basis, we could then ensure that TTL-related volumetric changes get properly "factored in." Unfortunately, we do not have TTL data available at this time.
- *Region-by-Region (or Sensor-by-Sensor) Analysis?* Currently we pool sensor traffic across all regions so that North American traffic is intermingled with European traffic, Asian traffic, African traffic, South American traffic, Australian/New Zealand/Pacific Ocean traffic, etc. The various regions have been and will be impacted by the virus differently, so it would be great if we could ask to see just traffic from North American sensors or just traffic from Asian sensors, for example. Doing a sensor-by-sensor analysis would also let is isolate/control for any changes that individual sensors may make to their configurations. Conceptually, this would potentially involve adding yet another item (sensor ID) to the current RRname, RRtype, Bailiwick, Rdata tuples. Unfortunately, we do not have sensor-tagged data available at this time.

Section IV: Acknowledgements

The author would like to thank:

- All of Farsight's sensor operators, for routinely contributing the data that this report relies upon
- Farsight's entire Engineering team, for their above-and-beyond effort in producing the software that makes Farsight's operations and these analyses possible
- Farsight's entire Ops crew, for providing stable systems and networks and for always answering inconveniently timed alerts when things get interesting due to no fault of their own
- Mr. Eric Ziegast, my colleague and fellow member of the Farsight Research Team, for his foresight and careful work in archiving and making available the daily files this analysis needed
- Ms. Karen Burke, Director of Corporate Communications at Farsight, for proposing this report in the first place, and for her careful shepherding and editing of it for eventual issuance
- Mr. Emanuel Younanzadeh and Ms. Rebecca Steiner of Farsight's Marketing Department, for a wonderful job with final editing, cover artwork, and production
- Mr. Ben April, Farsight's Chief Technology Officer, for his leadership, and for his careful review and numerous helpful suggestions for this report in particular

Any remaining errors or omissions are solely the responsibility of the author.

And let us close by wishing everyone the best of luck when it comes to getting safely through this terrible pandemic.

Appendix I. Timeline of Selected Coronavirus-Related Events

There have been many notable events that have already taken place during the COVID-19 pandemic. We know many more important events will unfold over time. Many factors complicate selection of events for inclusion in this timeline, including:

- The global extent of the pandemic
- The distributed nature of decision making in response to pandemic challenges, and
- The numerous different measures that have been undertaken in an effort to mitigate the impact of the disease

Nonetheless, we still wanted to ensure that we memorialized at least some benchmark event timing for context for our report. Inclusion (or failure to include) any event should not be taken to minimize or disregard the importance of any other particular event. We urge you to consult other timeline resources for more inclusive coverage, including:

- "Timeline of the COVID-19 pandemic"
  https://en.wikipedia.org/wiki/Timeline_of_the_COVID-19_pandemic

- "How the Coronavirus Pandemic Unfolded: a Timeline"
  https://www.nytimes.com/article/coronavirus-timeline.html"

With that for context, here's a condensed run down of some of the noteworthy events that have occurred during the pandemic:

- Dec 30th, 2019 "Unexplained pneumonia" in Wuhan, China's 10th largest city, noted in professional and public media (see https://promedmail.org/promed-post/?id=6864153 and https://www.reuters.com/article/us-china-health-pneumonia-idUSKBN1YZ0GP)

- Jan 13th, 2020 1st Case Outside China (this case was in Thailand)
  https://www.who.int/csr/don/14-january-2020-novel-coronavirus-thailand-ex-china/en/

- Jan 20th, 2020 Three distinct strains identified, indicative of the virus readily mutating
  http://weekly.chinacdc.cn/en/article/id/a3907201-f64f-4154-a19e-4253b453d10c

- Jan 21st, 2020 1st US Case (this was in Washington State)
  https://www.cdc.gov/media/releases/2020/p0121-novel-coronavirus-travel-case.html

- Jan 22nd, 2020 China Isolates Wuhan
  https://twitter.com/ChinaDaily/status/1220052882596286465

- Jan 24th, 2020 1st Three Cases in Europe (in France)
  https://solidarites-sante.gouv.fr/actualites/presse/communiques-de-presse/article/trois-cas-d-infection-par-le-coronavirus-2019-ncov-en-france-429100

- Jan 30th, 2020 WHO Declares Global Health Emergency
  https://www.nytimes.com/2020/01/30/health/coronavirus-world-health-organization.html

- Jan 31st, 2020 President Trump suspends entry of foreign nationals who had traveled to China in past 14 days
  https://www.whitehouse.gov/presidential-actions/proclamation-suspension-entry-immigrants-nonimmigrants-persons-pose-risk-transmitting-2019-novel-coronavirus/

- Feb 2nd, 2020 1st Death Outside China
  https://www.doh.gov.ph/press-release/DOH-reveals-more-negative-2019-nCoV-cases-confirms-first-nCoV-ARD-death-in-PH

- Feb 2nd, 2020 US Travel Restrictions Take Effect
  https://www.cnn.com/2020/02/02/us/coronavirus-us-travel-restrictions/index.html

- Feb 3rd, 2020 Diamond Princess cruise ship (accounting for more than half of all cases outside of China) quarantined in Yokohama, Japan
  https://www.theguardian.com/world/live/2020/feb/20/coronavirus-live-updates-diamond-princess-cruise-ship-japan-deaths-latest-news-china-infections

- Feb 5th, 2020 WHO Says "No Known Effective Treatment"
  https://www.reuters.com/article/us-china-health-treatments-who-idUSKBN1ZZ1M6

- Feb 11th, 2020 WHO officially names disease "COVID-19"
  https://twitter.com/WHO/status/1227248333871173632

- Feb 14th, 2020 First Case of Coronavirus in Africa
  https://www.aljazeera.com/news/2020/02/egypt-confirms-coronavirus-case-africa-200214190840134.html

- Feb 23rd, 2020 Major Outbreak in Italy
  https://www.nytimes.com/2020/02/23/world/europe/italy-coronavirus.html

- Feb 26th, 2020 CDC Confirms Possible Instance of Community Spread of COVID-19 in US
  https://www.cdc.gov/media/releases/2020/s0226-Covid-19-spread.html

- Feb 27th, 2020 Dow Falls 1,191 Points -- The Most in History
  https://www.cnn.com/2020/02/27/investing/dow-stock-market-selloff/index.html

- Feb 28th, 2020 Shelves emptied of toilet paper, hand sanitizer, bottled water and face masks as shoppers stock up
  https://www.usatoday.com/story/money/2020/02/28/coronavirus-2020-preparation-more-supply-shortages-expected/4903322002/

- Feb 29th, 2020 1st US Death (in Washington State)
  https://www.cdc.gov/media/releases/2020/s0229-COVID-19-first-death.html

- Mar 6th, 2020 CDC Says Those 60+ And With Underlying Medical Conditions Should Stay Home
  https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/people-at-higher-risk.html

- Mar 6th, 2020 Moscow Rejects Oil Production Cuts
  https://www.themoscowtimes.com/2020/03/06/russia-rejects-opec-saudi-arabias-coronavirus-oil-production-cut-a69553

- Mar 11th, 2020 Travel to US by most Visitors from Europe Blocked
  https://www.washingtonpost.com/world/2020/03/11/coronavirus-live-updates/

- Mar 11th, 2020 WHO Declares Coronavirus A Pandemic
  https://twitter.com/who/status/1237777021742338049?s=21

- Mar 11th, 2020 Colleges Ask Students to Leave Campuses
  https://www.insidehighered.com/news/2020/03/11/harvard-cornell-mit-and-others-ask-students-leave-campus-due-coronavirus

- Mar 13th, 2020 President Trump Declares National Emergency
  https://www.whitehouse.gov/presidential-actions/proclamation-declaring-national-emergency-concerning-novel-coronavirus-disease-covid-19-outbreak/

- Mar 15th, 2020 CDC Says No Gatherings of 50 or more
  https://www.npr.org/2020/03/15/816245252/cdc-recommends-suspending-gatherings-of-50-or-more-people-for-the-next-8-weeks

- Mar 16th, 2020 Statewide K12 School Closures Begin
  https://www.edweek.org/ew/section/multimedia/map-coronavirus-and-school-closures.html

- Mar 17th, 2020 EU Closes Borders to Most Non-Essential Travel
  https://www.cnn.com/2020/03/16/europe/spain-coronavirus-lockdown-intl/index.html

- Mar 23rd, 2020 WHO Says the Pandemic Is Accelerating
  https://www.bbc.com/news/world-52010304

- Mar 24th, 2020 Tokyo Olympics Delayed Until 2021
  https://www.npr.org/2020/03/24/820957235/tokyo-summer-olympics-postponed-to-2021

- Mar 26th, 2020 US $2 Trillion Stimulus Passed, Largest Aid Package in US History
  https://www.forbes.com/sites/jackbrewster/2020/03/27/trump-signs-2-trillion-stimulus-bill-into-law-largest-aid-package-in-us-history/#103b55d24ea5

- Mar 30th, 2020 316 million people in at least 42 states and elsewhere are now being urged to stay home
  https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html

- Apr 2nd, 2020 President Trump Uses Defense Production Act to order production of ventilators and protective masks
  https://www.whitehouse.gov/presidential-actions/memorandum-order-defense-production-act-regarding-purchase-ventilators/

- Apr 3rd, 2020 Justice Department authorizes use of home confinement for vulnerable inmates in Federal prisons
  https://www.fd.org/sites/default/files/covid19/bop_jail_policies_and_information/barr_memo_caresact_apr3_2020.pdf

- Apr 5th, 2020 President Trump Urges Use of Hydroxychloroquine As A Treament for Coronavirus
  https://www.nytimes.com/2020/04/05/us/politics/trump-hydroxychloroquine-coronavirus.html

- Apr 7th, 2020 President Trump criticizes WHO and Threatens to Pull US Funding
  https://www.npr.org/sections/coronavirus-live-updates/2020/04/07/829244345/trump-criticizes-who-and-threatens-to-pull-u-s-funding

- Apr 12th, 2020 First Stimulus Checks Deposited
  https://www.cbsnews.com/news/stimulus-checks-irs-deposits-first-wave-of-stimulus-checks-2020-04-12/

- Apr 13th, 2020 Reports of Potential Meat Shortages Due to Closure of Meat Processing Plants
  https://abcnews.go.com/Business/wireStory/virus-closes-meat-plants-ising-fears-shortages-70129905

- Apr 14th, 2020 IMF Say Experiencing Worst Global Economic Downturn Since the Depression
  https://blogs.imf.org/2020/04/14/the-great-lockdown-worst-economic-downturn-since-the-great-depression/

- Apr 14th, 2020 Top US General: "likely occurred naturally, as opposed to being created in a laboratory in China, but there is no certainty either way"
  https://www.reuters.com/article/us-health-coronavirus-usa-china-pentagon/u-s-military-says-coronavirus-likely-occurred-naturally-but-not-certain-idUSKCN21W2UH

- Apr 16th, 2020 President Trump releases roadmap for state reopening
  https://time.com/5822955/trump-unveils-phased-economy-reopening-coronavirus/

- Apr 20th, 2020 US, Canada and Mexico extend restrictions on non-essential cross-border travel
  https://www.dhs.gov/news/2020/04/20/acting-secretary-chad-wolf-statement-non-essential-travel

- Apr 20th, 2020 Oil Prices Go Negative
  https://www.npr.org/sections/coronavirus-live-updates/2020/04/20/838521862/free-falling-oil-prices-keep-diving-as-demand-disappears

- Apr 22nd, 2020 US Suspends Immigration Due to US Labor Market
  https://www.whitehouse.gov/presidential-actions/proclamation-suspending-entry-immigrants-present-risk-u-s-labor-market-economic-recovery-following-covid-19-outbreak/

- Apr 26th, 2020 CDC Announced 6 New Symptoms That May Be Sign of Coronavirus Infection
  https://www.usatoday.com/story/news/health/2020/04/26/coronavirus-symptoms-cdc-adds-six-new-symptoms-covid-19/3029438001/

- Apr 27th, 2020 WHO Says US Federal Coronavirus Plan is Clear and Science Based
  https://www.reuters.com/article/us-health-coronavirus-who-us/who-says-u-s-federal-coronavirus-plan-is-clear-and-science-based-idUSKCN2292EC

- Apr 28th, 2020 US Coronavirus Cases Top 1 Million
  https://www.wsj.com/articles/coronavirus-latest-news-04-28-2020-11588063873

Appendix II: Understanding "Typical" vs. "Atypical" Graphs

- As part of our feedback from internal reviewers, we came to understand that some readers might have a hard time noticing the same patterns that we saw (or their absence), so we wanted to add an appendix tackling that issue.

- To begin with, let me begin by emphasizing that this is NOT about calling out unusual increase amounts (e.g., we're NOT calling out sites that rise by 2X vs. sites that rise by 5X vs. sites that rise by 10X). We're interested in the **shape**, not the **scale**.

- We'll begin by considering a "typical" graph for most sectors (except for higher education), our friend *.forbes.com, as previously decomposed in the body of the report:

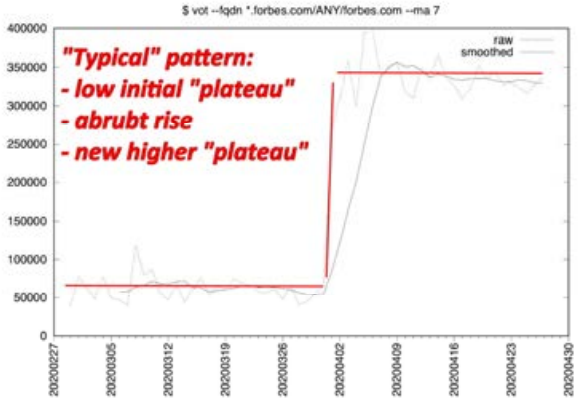

- When it comes to atypical graphs, there are many ways a graph can FAIL to fit those patterns, including:

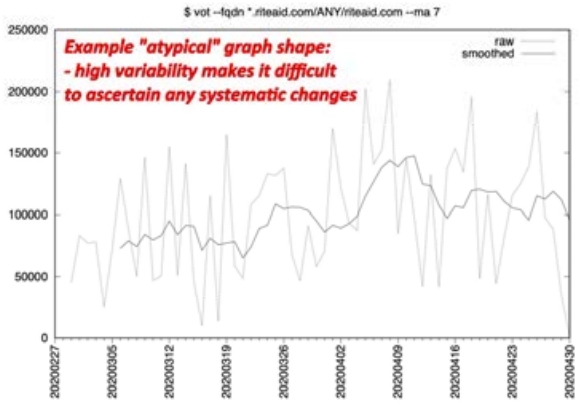

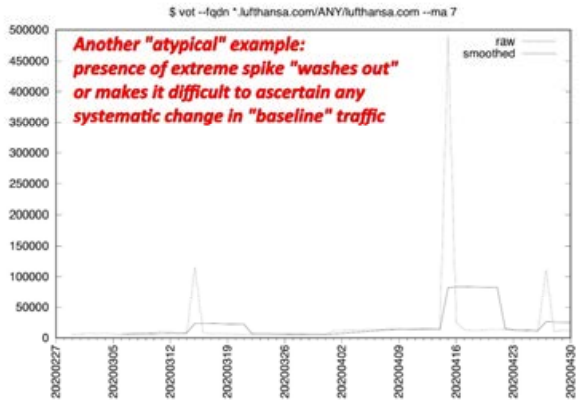

- When it comes to higher education, for reasons not yet understood, we saw a somewhat different characteristic graph shape:

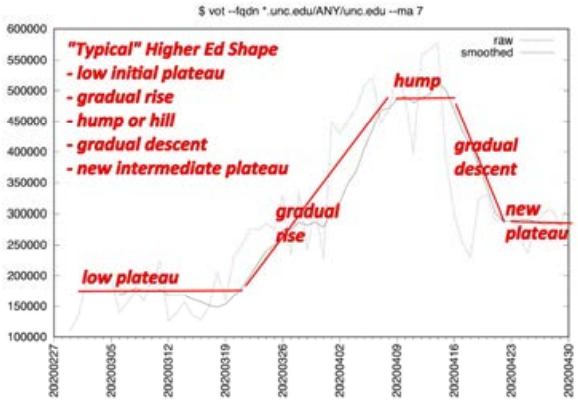

- Just like non-higher ed sites, higher ed sites had examples that didn't fit that pattern, including:

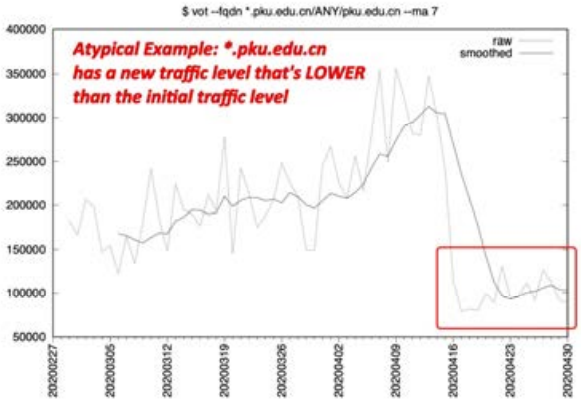

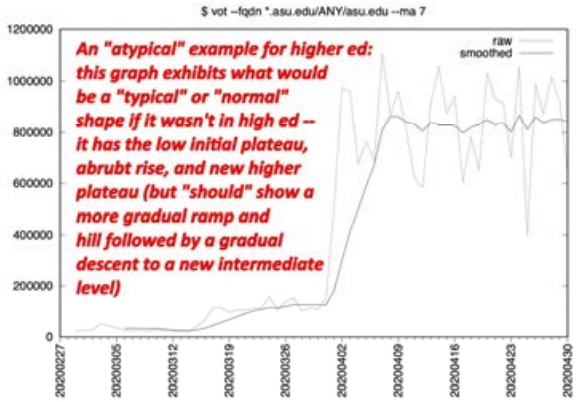

- Having offered this brief explanation-by-way-of-example, please don't get "hung up" on our typical/atypical characterization. While we're quite confident that some of these graphs are NOT the same as others, if you disagree with out tagging, feel free to disregard our tags and substitute whatever alternative works best for you (or use none at all).

FARSIGHT SECURITY

About Farsight Security

Farsight Security®, Inc. is the world's largest provider of historical and real-time DNS intelligence solutions. We enable security teams to qualify, enrich and correlate all sources of threat data and ultimately save time when it is most critical - during an attack or investigation. Our solutions provide enterprise, government and security industry personnel and platforms with unmatched global visibility, context and response. Farsight Security is headquartered in San Mateo, California, USA. Learn more about how we can empower your threat platform and security team with Farsight Security passive DNS solutions at https://www.farsightsecurity.com or follow us on Twitter: @FarsightSecInc.

To schedule a demo and obtain a free trial, contact: sales@farsightsecurity.com

+1-650-489-7919

Farsight Security®, Inc.177 Bovet Rd Ste 180 San Mateo, CA 94402 USA
info@farsightsecurity.com www.farsightsecurity.com