PonyStealer Infrastructure and Malware Analysis (Part 1/2)

Wes Connell Follow
Sep 25, 2018 · 7 min read

When I open-sourced StreamingPhish in late April 2018, I immediately set up an automated information gathering framework for each SSL certificate flagged with a score of 60% or higher (everything scored as “high”, “suspicious”, or “low”). This information includes metadata like name server records, ASNs, netblocks, IP owners, passive DNS, HTML, protocol headers, favicons, screenshots, and more (if you’re interested in this data, please reach out). I was spot-checking my tooling for taking webpage screenshots in early September and I noticed the following open directory at hxxps://accountinfologin[.]co.uk:

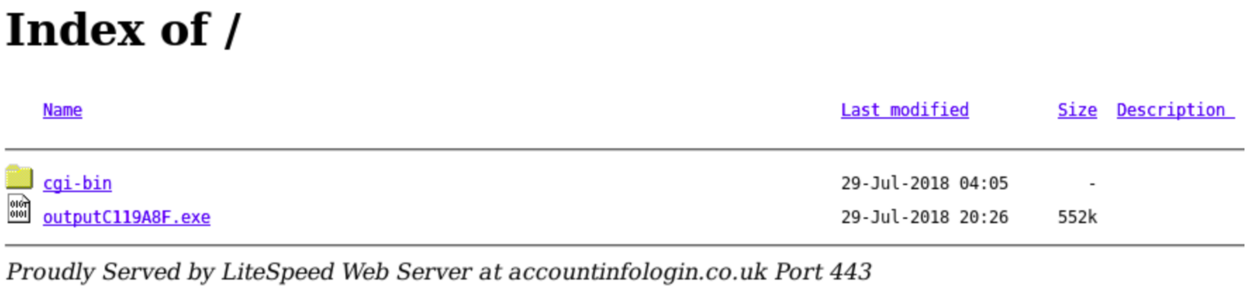

Index of /

Name	Last modified	Size	Description
cgi-bin	29-Jul-2018 04:05	-	
outputC119A8F.exe	29-Jul-2018 20:26	552k	

Proudly Served by LiteSpeed Web Server at accountinfologin.co.uk Port 443

Open directory at accountinfologin[.]co.uk.

The StreamingPhish utility is designed to detect phishing campaigns, so it’s far more common to find zip files of phishing kits in these open directories — not executable files. In this blog post, I’ll walk through my process for exploring the underlying infrastructure and discovering the intent of the executable file.

Infrastructure Analysis

A Certificate Authority (CA) can’t validate identities for domains that aren’t registered, so the first step is to take a look at the domain registration information. Using the “whois” utility from the command-line reveals the domain was registered on July 29th, 2018 with the registrar “1 & 1". GDPR constraints have forced organizations like ICANN to redact or completely omit useful information like the registrant’s name, e-mail address, or phone number, which might otherwise be valuable data points to explore additional domains registered by this actor:

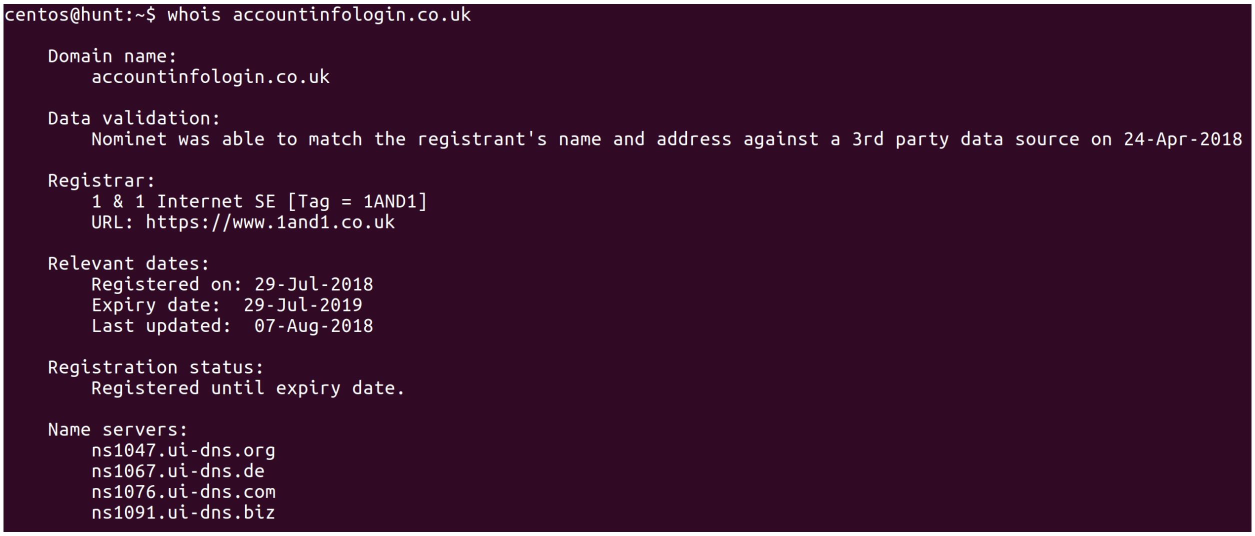

```
centos@hunt:~$ whois accountinfologin.co.uk

    Domain name:
        accountinfologin.co.uk

    Data validation:
        Nominet was able to match the registrant's name and address against a 3rd party data source on 24-Apr-2018

    Registrar:
        1 & 1 Internet SE [Tag = 1AND1]
        URL: https://www.1and1.co.uk

    Relevant dates:
        Registered on: 29-Jul-2018
        Expiry date:  29-Jul-2019
        Last updated:  07-Aug-2018

    Registration status:
        Registered until expiry date.

    Name servers:
        ns1047.ui-dns.org
        ns1067.ui-dns.de
        ns1076.ui-dns.com
        ns1091.ui-dns.biz
```

Registration information courtesy of Nominet.

The screenshot of the open directory showed the executable file was last modified on the same day the domain was registered, so it seems the open directory with that file was up for at least an entire month.

Next is to learn more about the hosting provider behind the web server. It's important that any DNS lookups for this step take place from a proxy or a virtual private server. If the actor behind the infrastructure is personally operating the authoritative name server, the IP address from your lookup could be observed the actor. Threat actors will personally operate these name servers for attacks like DNS tunneling, for example.

The "nslookup" utility from my virtual private server shows that this domain resolves to 162.244.94[.]135. I then used DNSdumpster to better understand which services might be running on the target. The output below indicates HTTP, HTTPS, and FTP services, including an adjacent server at 162.244.94[.]136:

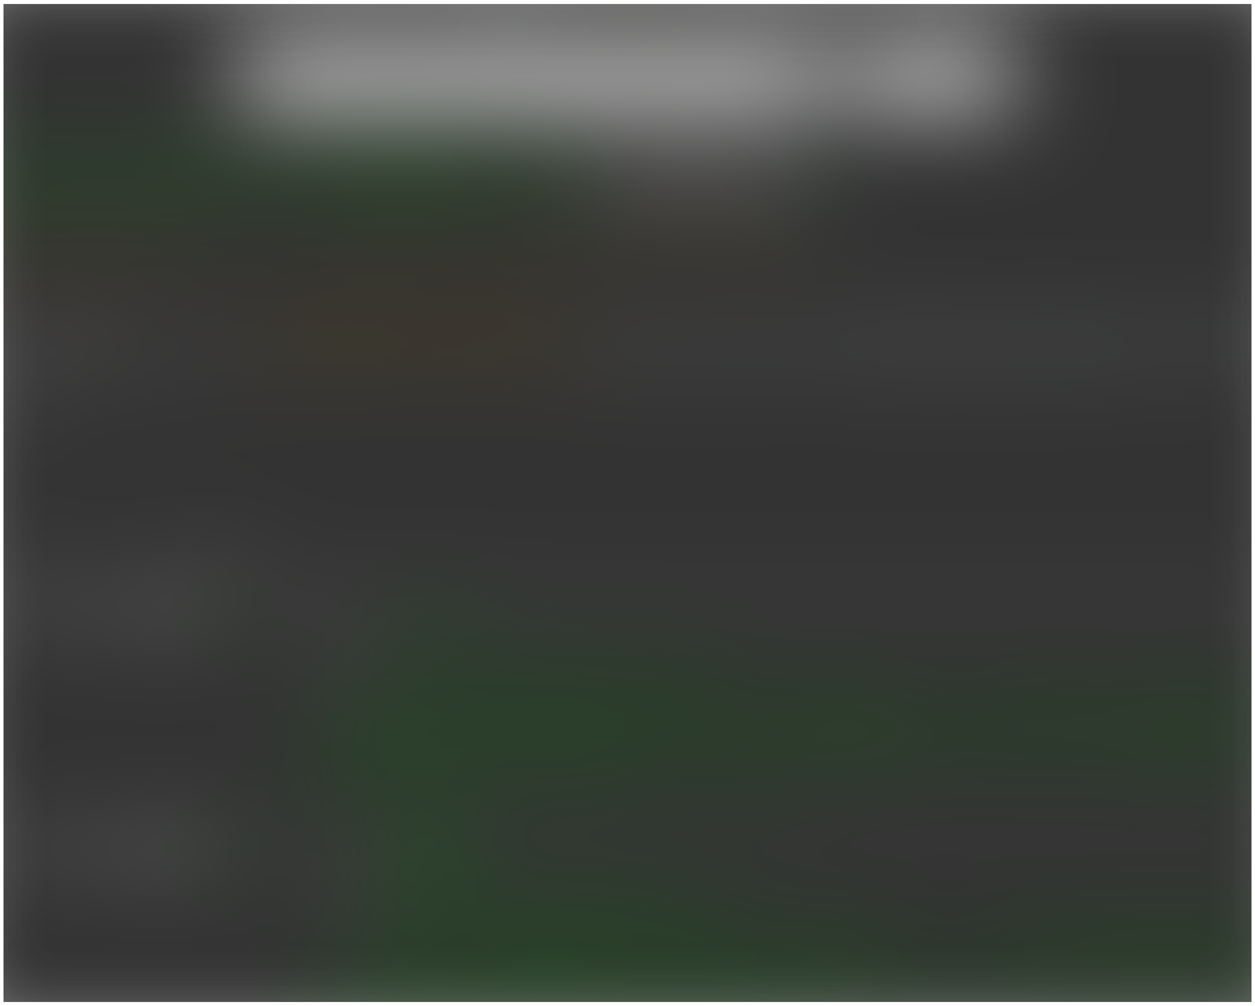

DNSdumpster results for the resolving IP address behind accountinfologin[.]co.uk.

Another useful tool for enumerating open ports and running services on an IP address is Shodan, which reveals a lot more than what was initially shown by DNSdumpster: IMAP, SMTP, cPanel, and MySQL just to name a few:

Shodan report for 162.244.94[.]135.

Next, I used the pyasn library in Python to map the IP address to it’s ASN (53667), and then used cidr-report to map the ASN to it’s hosting provider (PONYNET — FranTech Solutions, US):

ASN and hosting provider information behind 162.244.94[.]135.

I’m not sure of FranTech’s reputation off the top of my head, but I do remember Brian Kreb’s detailing an occasion last year where the owner of FranTech ignored abuse reports of IoT botnet control channels it was hosting, and subsequently got slammed by the Mirai botnot, knocking FranTech's infrastructure completely offline. Querying my StreamingPhish data collector for SSL certificates whose Common Name (CN) resolves to an IP address hosted by FranTech shows no shortage of very suspicious, high-scoring domains:

Infrastructure hosted by FranTech that's been flagged by StreamingPhish over the last few months.

Moving to passive DNS, PassiveTotal is my go-to service to explore historical domain name resolutions, and a free account gets you 15 free queries per day. There appears to be another 60 domains at first glance with lexical similarities to accountinfologin[.]co.uk that also resolve to 162.244.94[.]135, suggesting they could be part of the same campaign. The heatmap below indicates that these IP address resolutions began on July 28th and are still resolving to this day:

Heatmap from PassiveTotal of domain names resolving to 162.244.94[.]135, beginning in late July.

And here’s a snippet of the 60 domain names resolving to the IP:

Also from PassiveTotal. Note the lexical similarities that each of these domains have in common… ”acc”, ”info”, ”login”, all belonging to the co.uk top-level domain.

The list of 60 domain names enticed me to check and see if StreamingPhish caught any SSL certificates issued to them as well. It turns out that it detected 4 SSL certificates issued to these domains on August 31st, the same day accountinfologin[.]co.uk was issued it’s certificate:

Querying for the domains found in PassiveTotal against my StreamingPhish data collector yielded four hits (in addition to accountinfologin[.]co.uk).

I was surprised to see so few certificates detected by StreamingPhish — just 4 of potentially 60. Was it because SSL certificates weren't issued to these domains, or because my classifier was missing them? Unfortunately, it seems to be a little bit of both. The existing feature vector in the classifier wasn't identifying enough phish-y characteristics to warrant a high score:

```
  [Phishing]      accountinfologin.co.uk  0.827
  [Phishing]      accountinfologinservices.co.uk  0.953
  [Phishing]      www.accountinfologinservices.co.uk      0.953
  [Not Phishing]  acc-info-login.co.uk    0.367
  [Not Phishing]  accinfologinshop.co.uk  0.091
  [Not Phishing]  accinfoservicessolutions.co.uk  0.007
  [Not Phishing]  www.accinfoservicessolutions.co.uk      0.007
  [Not Phishing]  theaccinfologin.co.uk   0.127
  [Not Phishing]  accinfologin.co.uk      0.089
  [Not Phishing]  accinfologinservices.co.uk      0.304
  [Not Phishing]  acc-infos-logined.co.uk 0.177
  [Not Phishing]  theaccinfoslogin.co.uk  0.140
  [Not Phishing]  myaccinfologin.co.uk    0.080
```

```
[Phishing]      theaccountinfologin.co.uk        0.859
[Not Phishing]  accinfoslogined.co.uk    0.127
[Not Phishing]  theaccinfoslogined.co.uk 0.152
[Not Phishing]  acc-info-services.co.uk  0.016
[Phishing]      account-info-login.co.uk 0.997
[Not Phishing]  poawrizr.co.uk   0.000
[Phishing]      www.theaccountinfologin.co.uk    0.859
[Not Phishing]  www.accinfologin.co.uk   0.089
[Not Phishing]  www.acc-info-login.co.uk 0.367
[Not Phishing]  accinfoslogin.co.uk      0.102
[Not Phishing]  www.accinfoslogin.co.uk  0.102
[Not Phishing]  www.acc-infos-login.co.uk        0.478
[Not Phishing]  acc-infos-login.co.uk    0.478
[Not Phishing]  www.theaccinfoservices.co.uk     0.007
[Not Phishing]  theaccinfoservices.co.uk 0.007
[Not Phishing]  www.theaccinfologin.co.uk        0.127
[Not Phishing]  accinfoservicesonline.co.uk      0.043
[Not Phishing]  www.accinfoservicesonline.co.uk  0.043
[Not Phishing]  www.theaccinfoslogin.co.uk       0.140
[Not Phishing]  www.myaccinfologin.co.uk 0.080
[Not Phishing]  www.accinfoservices.co.uk        0.006
[Not Phishing]  www.accinfologinservices.co.uk   0.304
[Not Phishing]  accinfoservices.co.uk    0.006
[Not Phishing]  www.accinfoslogined.co.uk        0.127
[Not Phishing]  www.poawrizr.co.uk       0.000
[Not Phishing]  www.theaccinfoslogined.co.uk     0.152
[Not Phishing]  www.acc-infos-logined.co.uk      0.177
[Not Phishing]  www.accinfoservicesexpress.co.uk 0.006
[Not Phishing]  accinfoservicesexpress.co.uk     0.006
[Not Phishing]  accinfosloginedservices.co.uk    0.318
[Not Phishing]  www.accinfosloginedservices.co.uk0.318
[Not Phishing]  accinfosloginservices.co.uk      0.301
[Not Phishing]  www.accinfosloginservices.co.uk  0.301
[Phishing]      www.myaccountinfologin.co.uk     0.971
[Not Phishing]  www.myaccinfoservices.co.uk      0.005
[Not Phishing]  www.accinfosloginshop.co.uk      0.091
[Phishing]      myaccountinfologin.co.uk 0.971
[Not Phishing]  accinfosloginshop.co.uk  0.091
[Not Phishing]  myaccinfoservices.co.uk  0.005
[Not Phishing]  www.myaccinfoslogined.co.uk      0.108
[Not Phishing]  myaccinfoslogined.co.uk  0.108
[Not Phishing]  myaccinfoslogin.co.uk    0.089
[Not Phishing]  www.myaccinfoslogin.co.uk        0.089
[Not Phishing]  www.acc-info-services.co.uk      0.016
[Not Phishing]  www.accinfologinshop.co.uk       0.091
[Phishing]      www.account-info-login.co.uk     0.997
```

```
[Phishing]      www.accountinfologin.co.uk       0.827
```

I added three keyword matches to my feature vector (“acc”, “info”, and “login”), did a retrain using the existing training data, and fortunately the classifier now detects all but two of the domains.

```
[Phishing]      accountinfologin.co.uk   0.997
[Phishing]      accountinfologinservices.co.uk   1.000
[Phishing]      www.accountinfologinservices.co.uk       1.000
[Phishing]      acc-info-login.co.uk     0.997
[Phishing]      accinfologinshop.co.uk   0.992
[Phishing]      accinfoservicessolutions.co.uk   0.841
[Phishing]      www.accinfoservicessolutions.co.uk       0.841
[Phishing]      theaccinfologin.co.uk    0.990
[Phishing]      accinfologin.co.uk       0.985
[Phishing]      accinfologinservices.co.uk       0.998
[Phishing]      acc-infos-logined.co.uk  0.999
[Phishing]      theaccinfoslogin.co.uk   0.991
[Phishing]      myaccinfologin.co.uk     0.987
[Phishing]      theaccountinfologin.co.uk        0.998
[Phishing]      accinfoslogined.co.uk    0.990
[Phishing]      theaccinfoslogined.co.uk 0.992
[Phishing]      acc-info-services.co.uk  0.820
[Phishing]      account-info-login.co.uk 1.000
[Not Phishing]  poawrizr.co.uk   0.005
[Phishing]      www.theaccountinfologin.co.uk    0.998
[Phishing]      www.accinfologin.co.uk   0.985
[Phishing]      www.acc-info-login.co.uk 0.997
[Phishing]      accinfoslogin.co.uk      0.987
[Phishing]      www.accinfoslogin.co.uk  0.987
[Phishing]      www.acc-infos-login.co.uk        1.000
[Phishing]      acc-infos-login.co.uk    1.000
[Phishing]      www.theaccinfoservices.co.uk     0.833
[Phishing]      theaccinfoservices.co.uk 0.833
[Phishing]      www.theaccinfologin.co.uk        0.990
[Phishing]      accinfoservicesonline.co.uk      0.966
[Phishing]      www.accinfoservicesonline.co.uk  0.966
[Phishing]      www.theaccinfoslogin.co.uk       0.991
[Phishing]      www.myaccinfologin.co.uk 0.987
```

```
[Phishing]      www.accinfoservices.co.uk        0.805
[Phishing]      www.accinfologinservices.co.uk   0.998
[Phishing]      accinfoservices.co.uk    0.805
[Phishing]      www.accinfoslogined.co.uk        0.990
[Not Phishing]  www.poawrizr.co.uk       0.005
[Phishing]      www.theaccinfoslogined.co.uk     0.992
[Phishing]      www.acc-infos-logined.co.uk      0.999
[Phishing]      www.accinfoservicesexpress.co.uk 0.823
[Phishing]      accinfoservicesexpress.co.uk     0.823
[Phishing]      accinfosloginedservices.co.uk    0.998
[Phishing]      www.accinfosloginedservices.co.uk0.998
[Phishing]      accinfosloginservices.co.uk      0.998
[Phishing]      www.accinfosloginservices.co.uk  0.998
[Phishing]      www.myaccountinfologin.co.uk     1.000
[Phishing]      www.myaccinfoservices.co.uk      0.824
[Phishing]      www.accinfosloginshop.co.uk      0.992
[Phishing]      myaccountinfologin.co.uk 1.000
[Phishing]      accinfosloginshop.co.uk  0.992
[Phishing]      myaccinfoservices.co.uk  0.824
[Phishing]      www.myaccinfoslogined.co.uk      0.991
[Phishing]      myaccinfoslogined.co.uk  0.991
[Phishing]      myaccinfoslogin.co.uk    0.989
[Phishing]      www.myaccinfoslogin.co.uk        0.989
[Phishing]      www.acc-info-services.co.uk      0.820
[Phishing]      www.accinfologinshop.co.uk       0.992
[Phishing]      www.account-info-login.co.uk     1.000
[Phishing]      www.accountinfologin.co.uk       0.997
```

I'd get even higher scores if I extended my malicious training set to include the domains I'm evaluating (but that's akin to having answers to an exam before you take it, then bragging about getting an "A"). I'll plan to update the training data in StreamingPhish to include these domains in a later release, however.

That covers most of the infrastructure analysis for now. In a future post, I'll walk through my approach for understanding the intent of the executable file in that open directory (static analysis, dynamic analysis, interactive behavioral analysis, and code-level disassembly). I’ve only downloaded the file thus far and run it’s SHA256 hash against VirusTotal. Several of the engines marked it as malicious and indicated it’s a variant of PonyStealer, a well-known information-stealing piece of malware:

SHA256 lookup from outputC119A8F.exe, which was sitting in an open directory at accountinfologin[.]co.uk.

Thanks for reading!