

Drawing a Dragon: Connecting the Dots to Find APT41

 blogs.blackberry.com/en/2021/10/drawing-a-dragon-connecting-the-dots-to-find-apt41

The BlackBerry Research & Intelligence Team



Executive Summary

The [BlackBerry Research & Intelligence Team](#) recently connected seemingly disparate malware campaigns, which began with an unusual Cobalt Strike configuration that was first included in a [blog post](#) published the same month as COVID-19 lockdowns began in Europe and the U.S. What we found led us through a malicious infrastructure that had been partially documented in articles by several other research organizations.

The image we uncovered was that of a state-sponsored campaign that plays on people's hopes for a swift end to the pandemic as a lure to entrap its victims. And once on a user's machine, the threat blends into the digital woodwork by using its own customized profile to hide its network traffic.

Introduction

APT41 is a prolific Chinese state-sponsored cyberthreat group that has conducted malware campaigns related to espionage and financially motivated criminal activity dating as far back as [2012](#). This threat group has targeted organizations around the world, in verticals such as travel, telecommunications, healthcare, news and education.

APT41 has often used [phishing emails](#) with malicious attachments as an initial infection vector. Once it has gained access to a target organization, it typically deploys more advanced malware to establish a persistent foothold. This group uses a variety of different malware families including information stealers, keyloggers and backdoors.

BlackBerry researchers have been monitoring Cobalt Strike activity that used a bespoke, malleable command-and-control (C2) profile, which had settings that were previously documented in a report by FireEye in March of 2020. They attributed this configuration to APT41-related activity.

We were able to uncover what we believe is additional APT41 infrastructure by taking these unique aspects and following the trail of digital breadcrumbs. Overlapping indicators of compromise (IOCs) linked the trail of our findings to those of two additional campaigns documented by [Positive Technologies](#) and [Prevailion](#). These posts were titled "*Higaisa or Winnti? APT41 backdoors, old and new,*" and "*The Gh0st Remains the Same,*" respectively.

We also found three additional phishing lures targeting victims in India, containing information related to new tax legislation and COVID-19 statistics. These messages masqueraded as being from Indian government entities.

These lures were part of an execution chain that had the goal of loading and executing a Cobalt Strike Beacon on a victim's network. The phishing lures and attachments also fit tactics that were previously used in infection vectors by APT41. These findings show that the APT41 group is still regularly conducting new campaigns, and that they will likely continue to do so in the future.

Connecting the Dots

A recent [blog post published by FireEye](#) in March of 2020 explored APT41's tactics, including their use of malicious documents, exploits and Cobalt Strike. The report indicated that the group was using a bespoke, malleable C2 profile with at least one of its Cobalt Strike Beacons.

A malleable C2 profile is a feature within Cobalt Strike that allows an attacker to customize a Beacon's network communications to its C2 channel in a way that allows it to blend into normal traffic on a victim network. For example, there are publicly available profiles that are designed to look like legitimate network traffic from Amazon, Gmail, OneDrive and [many others](#).

We uncovered a malleable C2 [profile on GitHub](#) that is very similar to that of the one mentioned in the FireEye blog. This one seems to have been authored by a Chinese security researcher with the pseudonym “1135.”

These profiles had several similarities: Both used jQuery Malleable C2 profiles, and portions of the HTTP GET profile block are almost identical. HTTP header fields such as “accept,” “user-agent,” “host,” and “referer,” as well as the “set-uri” field, were all exact matches to the profile data listed in the FireEye blog.

```
http-get {
  set uri "/jquery-3.3.1.min.js";
  set verb "GET";
  client {
    header "Accept" "text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8";
    header "Host" "cdn.bootcss.com";
    header "Referer" "http://cdn.bootcss.com/";
    header "Accept-Encoding" "gzip, deflate";
    metadata {
      base64url;
      prepend "__cfduid=";
      header "Cookie";
    }
  }
}
```

Figure 1: JQuery Malleable C2 from the ‘1135’ Github

Armed with this data point, we can perform some deeper visual analysis of Beacon configuration data in our possession to reveal patterns that are only perceptible when a large data set is accessible. By extracting and correlating the HTTP headers used in the GET and POST requests defined in the Beacon configs, we can generate revealing connections between seemingly disparate Cobalt Strike infrastructure.

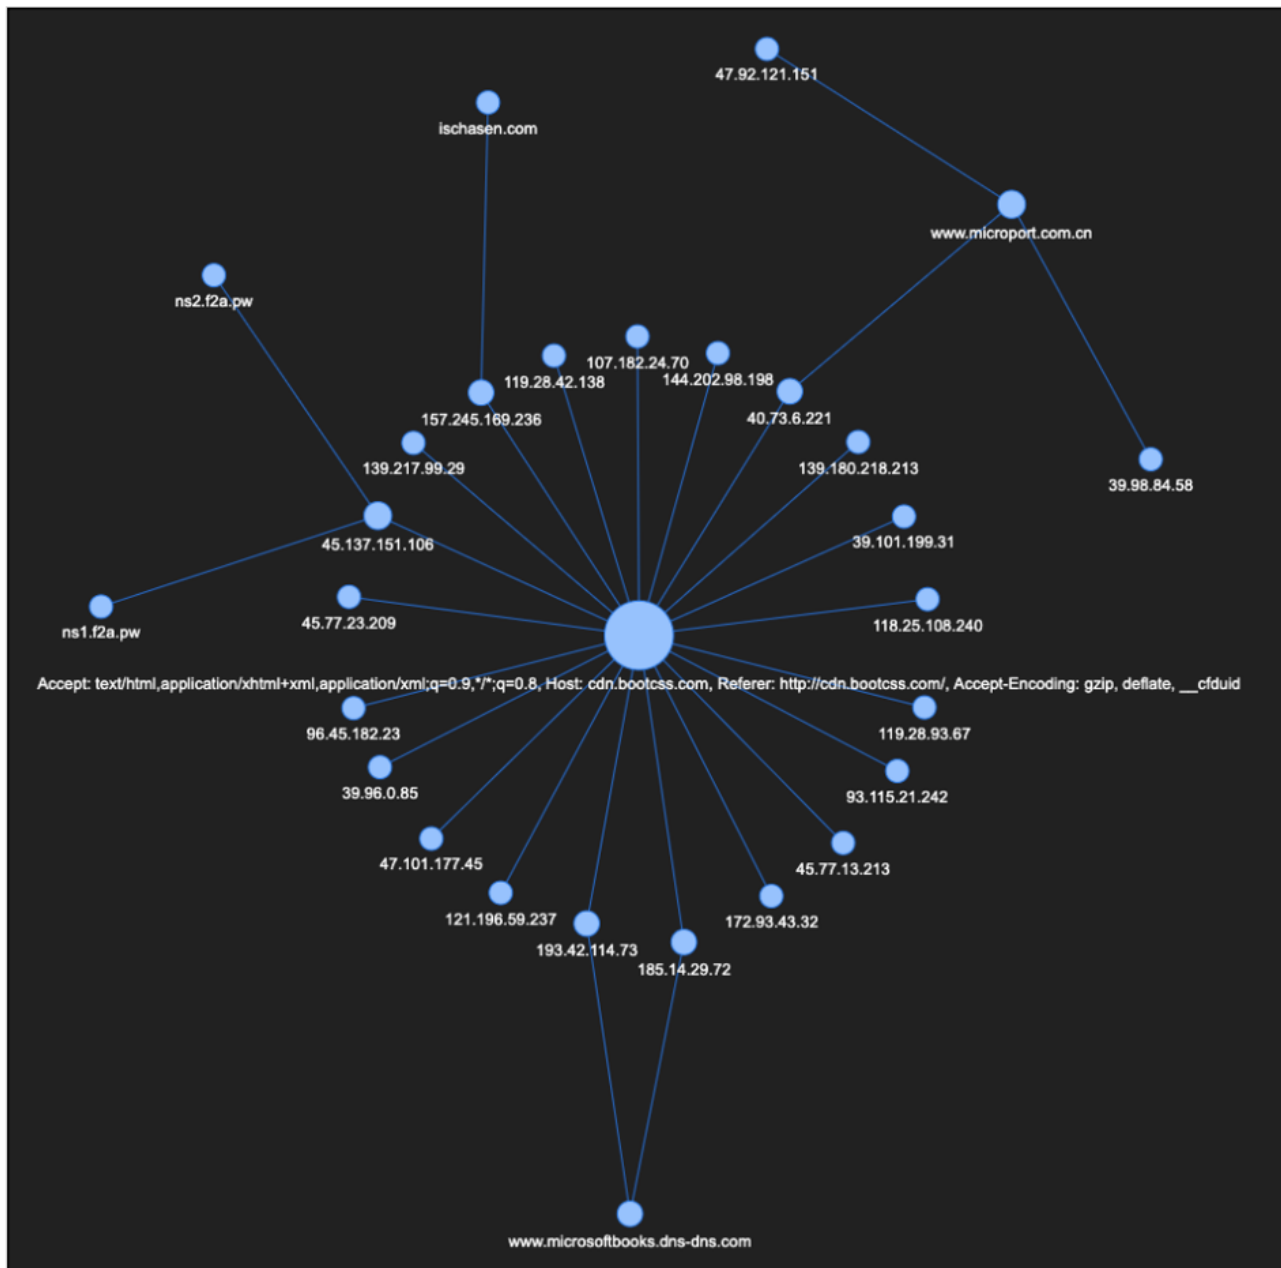


Figure 2: Clustering on "cdn.bootcss.com" HttpPost_Metadata

While we identified a relatively small number of Beacons using the BootCSS domain as part of their Malleable C2 configuration, there were also a few clusters with unique configuration metadata that enabled us to identify additional beacons related to APT41.

The Beacons served by these new nodes are using a different malleable profile to those in the original cluster that attempts to make the Beacon traffic look like legitimate Microsoft traffic.

IP	Domain
----	--------

144.202.98.198	zalofilescdn[.]com
-----------------------	--------------------

107.182.24.70	isbigfish[.]xyz
----------------------	-----------------

185.14.29.72	www[.]microsoftbooks[.]dns-dns[.]com
---------------------	--------------------------------------

193.42.114.73	www[.]microsoftbooks.dns-dns[.]com
----------------------	------------------------------------

149.28.78.89	www[.]mlcrosoft[.]site
---------------------	------------------------

23.67.95.153	ns[.]mircosoftdoc[.]com
---------------------	-------------------------

104.27.132.211	cdn[.]microsoftdocs[.]workers[.]dev
	ccdn[.]microsoftdocs[.]workers[.]dev

The domains we found share similarities in their naming convention, which try to masquerade as legitimate Microsoft® domains. Searching for these IPs and domains in a variety of open source intelligence tool (OSINT) repositories reveals some connections that bear further examination. The IP 107.182.24[.]70 as well as the domain www[.]mlcrosoft[.]site both appear within a blog from Positive Technologies. Further hunting for the IP address 149.28.78[.]89 reveals links to a campaign mentioned in the previously referenced Prevailion blog.

Gh0st in the Machine

In that blog, we can find two IOCs that appear in the cluster above; the IP 149.28.78.89, and the domain *mlcrosoft.site*. The blog associated those IOCs with the Higaisa advanced persistent threat (APT) group, which operates out of North Korea.

The domain mlcrosoft[.]site also appears in the blog from Positive Technologies. That article has additional overlapping IOCs, and talks about the same campaign as mentioned in the Prevailion blog. However, it makes a strong argument that the activity is from APT41 rather than Higaisa APT.

When we do a side-by-side comparison of the domains from the Positive Technologies blog and our datasets, there is a strong similarity between naming conventions used:

BlackBerry IOCs

Positive Technologies IOCs

www[.]microsoftbooks.dns-dns[.]com	microsoftbooks[.]dynamic-dns[.]net
cdn[.]microsoftdocs.workers[.]dev	microsoftdocs[.]dns05[.]com
ccdn[.]microsoftdocs.workers[.]dev	ns[.]microsoftdocs.dns05[.]com
ns[.]mircosoftdoc[.]com	ns1[.]microsoftsonline[.]net

We also discovered that *mlcrosoft[.]site* and *mircosoftdoc[.]com* both appear in the Azure-Sentinel detection rule for known Barium phishing domains. The IP 144.202.98[.]198 has also been previously associated with APT41/Barium by a Microsoft researcher.

Another IP from this cluster, 185.14.29[.]72, was recently providing virtual hosting for several domain names such as:

- chaindefend[.]bid
- defendchain.[.]xyz
- assistcustody[.]xyz
- microsoftonlineupdate.dynamic-dns[.]net

Previously, this IP has been associated with DNS resolutions for schememicrosoft[.]com and www.microsoftbooks[.]dns-dns.com. Several of the domains also have links to 209.99.40[.]222, an IP that is known to perform malicious DNS/bulletproof hosting.

As of Sept. 14, 2021, this IP resolved to a new domain very briefly: www.microsoftonlineupdate.dynamic-dns[.]net. This domain also conforms to a naming convention similar to those we have seen in the previous table.




Figure 3: Teamserver geolocation

Phishing Lures

By performing further intelligence correlation to investigate these URLs further, we found a malicious PDF that reaches out to `ccdn[.]microsoftdocs.workers[.]dev`. This site had previously hosted a Cobalt Strike Team Server.

Further digging reveals a set of three PDFs used as malicious phishing lures, which are linked to the `*.microsoftdocs.workers[.]dev` domains. These lures all target victims in India, either promising information regarding new India-specific income taxation rules or COVID-19 advisories.

The first lure – “Income tax new rules for NRI.pdf.lnk” – contains both a PDF document and an embedded PowerShell script. Upon execution, the PDF is displayed to the user, after which the PowerShell is executed in the background.

Income tax new rules for NRI: Everything on residential status and taxation in India

The COVID-19 crisis has brought about a steep increase in job layoffs across the world which has in turn had a huge impact on the Indian migrants residing abroad. A lot of Non-Resident Indians (NRIs), therefore, are looking to return to their hometown. If you are considering returning to India, a vigilant planning would be required for a hassle-free relocation. A generic concern amongst most NRIs is the tax implications on the income earned during the year in which NRI returns to India.

The Union Budget 2020 has reformulated the norms for determining the tax residential status of an individual in India. An individual in India can be classified into three categories - a) Resident and Ordinarily Resident (ROR); b) Resident but Not Ordinarily Resident (RNOR); and c) Non-Resident (NR).

Rules to determine residential status of NRI

Residential status of an individual is determined based on his/her physical presence in India during a Financial Year (FY). The Hon'ble Finance Minister, in Budget 2020 has brought about certain amendments to the residency provisions. Below is a glimpse of the amendments that one needs to keep in mind for determination of residential status.

If an individual satisfies any of the following two conditions he will qualify as a resident of India else will qualify as a NR for that FY:

- a) his/her stay exceeds 182 days during the FY or
- b) his/her stay exceeds 60 days during a FY and 365 days in preceding 4 FY's.

If an Indian citizen or Person of Indian Origin (PIO) comes for a visit to India, then 60 days mentioned in point b) above will be enhanced to 182 days.

There is a new leg to this condition, which states that if an Indian citizen or PIO who comes for a visit to India and his total income from a business/or a profession set up in India exceeds ₹15 lakh during that FY, then 60 days will substitute by 120 days. However, if the ₹15 lakh threshold is not met, then the 182 condition will apply.

Figure 4: Phishing lure 1

The PowerShell script downloads and executes a payload via “%temp%\conhost.exe,” which loads a payload file called “event.dat.” This .DAT file is a Cobalt Strike Beacon.

The second and third lures each have similar execution flows and component parts; a PDF lure, conhost.exe, and an event.* payload. In this case, these event files had a .LOG extension, rather than .DAT.

The biggest difference between the second and third lures is that one uses a self-extracting archive named “India records highest ever single day covid_19 recoveries.pdf.exe,” and the other uses a ZIP file named “India records highest ever single day COVID-19 recoveries.zip.”

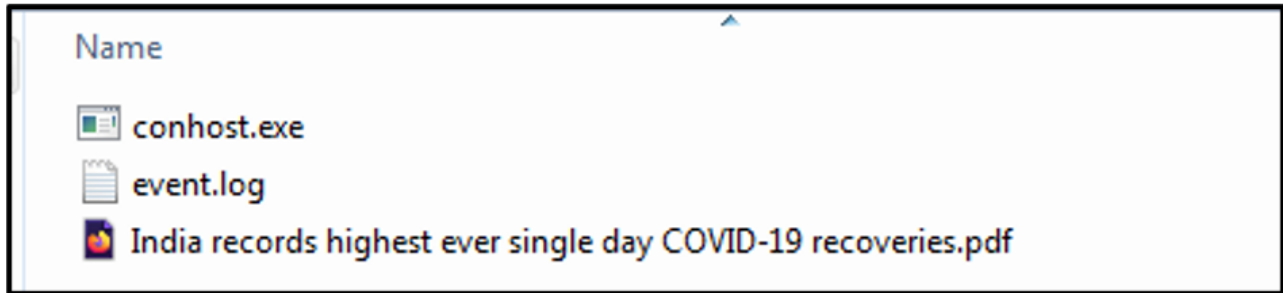


Figure 5: Contents of lures 2 and 3

Lures two and three also contain the same information within their respective PDFs. Both relate to a record high number of COVID-19 recoveries in India, information which purports to be from the Indian Government Ministry of Health & Family Welfare.

India records highest ever single day COVID-19 recoveries of 51,255 cases



Ministry of Health & Family Welfare

Government of India

A record 51,255 recoveries have been registered in the country in the last 24 hours. This is the highest number of recoveries in a single day since the outbreak of the pandemic in the country. With this, the recovery rate has improved to 65.43 per cent. The case fatality rate has further declined to 2.13 per cent.

The Health and Family Welfare Ministry said, a total of 11,45,629 people have recovered from COVID-19 in the country so far. 54,735 new cases of Coronavirus have been reported in the last 24 hours, taking the total number of cases to 17,50,723. Presently, the total number of active Corona cases in the country is 5,67,730. In a single day, 853 deaths have also been reported taking the nationwide toll to 37,364.

The Indian Council of Medical Research said that a total of 4,63,172 tests were conducted by various laboratories within 24 hours. Overall, 1,98,21,831 tests have been conducted in the country so far. At present, 1,344 laboratories across India are conducting COVID-19 tests that includes 913 government laboratories and 431 private laboratory chains.

Figure 6: Phishing lure 2 and 3 PDF contents

By extracting the configurations from each of the three lures' event.* Beacon payloads, we can see that the C2 server address used in the configuration data differs slightly:

- Lure-1 uses `ccdn[.]microsoftdocs.workers.dev/en-us/windows/apps/`
- Lure-2 & 3 use `cdn[.]microsoftdocs.workers.dev/en-us/windows/apps/`

The same can be seen for the `HttpGet_Metadata` and `HttpPost_Metadata` host addresses:

- Lure-1 uses `ccdn[.]microsoftdocs.workers.dev`
- Lure-2 & 3 use `cdn[.]microsoftdocs.workers.dev`

KEY	VALUE
BeaconType	HTTPS
Port	443
SleepTime	1000
MaxGetSize	1398104
Jitter	0
MaxDNS	255
C2Server	<ul style="list-style-type: none"> • ccdn[.]microsoftdocs.workers.dev,/en-us/windows/apps/ (Lure 1) • cdn[.]microsoftdocs.workers.dev,/en-us/windows/apps/ (Lure 2&3)
UserAgent	Mozilla/5.0 (MSIE 10; Windows NT 6.1; Trident/5.0)
HttpPostUri	/en-us/windows/windows-server/
Malleable_C2_Instructions	Base64 decode
HttpGet_Metadata	<p>ConstHeaders</p> <ul style="list-style-type: none"> • Host: ccdn[.]microsoftdocs.workers.dev (Lure-1) • Host: cdn[.]microsoftdocs.workers.dev(Lure-2&3) • User-Agent: Mozilla/5.0 (Windows NT 6.1; Trident/5.0) • Accept: */* • Accept-Encoding: gzip, deflate, br <p>Metadata</p> <ul style="list-style-type: none"> • base64 • prepend "__cfduid=" • header "Cookie"

HttpPost_Metadata	<p>ConstHeaders</p> <ul style="list-style-type: none"> • Host: ccdn[.]microsoftdocs.workers.dev (Lure-1) • Host: cdn[.]microsoftdocs.workers.dev(Lure-2&3) • User-Agent: Mozilla/5.0 (Windows NT 6.1; Trident/5.0) • Accept: */* • Accept-Encoding: gzip, deflate, br <p>SessionId</p> <ul style="list-style-type: none"> • base64url • parameter "k" <p>Output</p> <ul style="list-style-type: none"> • base64 • print
DNS_Idle	104.88.34.55
DNS_Sleep	0
HttpGet_Verb	GET
HttpPost_Verb	POST
HttpPostChunk	0
Spawnto_x86	%windir%\syswow64\gpupdate.exe
Spawnto_x64	%windir%\sysnative\gpupdate.exe
CryptoScheme	0
Proxy_Behavior	Use IE settings
Watermark	305419896
bStageCleanup	False

bCFGCaution	False
KillDate	0
bProInjct_StartRWX	True
bProInjct_UseRWX	True
bProInjct_MinAllocSize	0
ProInjct_PrependAppend_x86	Empty
ProInjct_PrependAppend_x64	Empty
ProInjct_Execute	<ul style="list-style-type: none"> • CreateThread • SetThreadContext • CreateRemoteThread • RtlCreateUserThread
ProInjct_AllocationMethod	VirtualAllocEx
bUsesCookies	True

A report by Subex from September 2020 described a campaign using similar phishing lures, also targeting Indian nationals, which they attributed to the [Evilnum APT group](#). The indicators of compromise (IOCs) in this report are quite similar (or even identical) to those of the phishing lures we've previously investigated. We believe that this attack was perpetrated by APT41 and not the Evilnum group for several reasons:

The first is that the *Event.** payloads are in fact Cobalt Strike Beacons, as per the extracted configuration data shown in the table above. This behavior is indicative of APT41 rather than the Golden Chickens Malware-as-a-Service (MaaS), as reported in the Subex Evilnum APT report.

The second reason is that there are several configuration settings that indicate APT41 activity when they're aggregated. These same settings were present within the phishing lure Beacons and have been observed in previous attacks by this group.

In addition to this, the aforementioned [blog from Positive Technologies](#) contained overlapping infrastructure that ties in with what we have observed. They documented similar phishing lures using PDF documents as bait, which they attributed to APT41.

These lures follow a similar naming convention to the ones we've documented here. Their execution chains encompass both a loader and payload component, but Cobalt Strike was just one of several potential payloads that were listed.

The use of spear-phishing attachments to gain initial access has been a known [APT41 tool, technique and procedure](#) (TTP) for years. In addition, the previously discussed overlap in network infrastructure adds credence to this being an APT41-affiliated campaign.

Conclusions

It's a rare treat for a research organization to have a truly robust set of data about any one threat. Having a security industry that puts a strong emphasis on the public sharing of information means that we can put our collective heads together to create a more complete picture.

An article posted by FireEye initially pointed us to a Malleable C2 profile for Cobalt Strike. We researched this and found a similar profile that had Beacons using the BootCSS domain as part of their configuration. This then pointed us to additional overlapping configuration metadata within the Beacon configuration, which subsequently steered us into identifying a whole new cluster and a new set of domains.

This discovery led us to find connections with the campaign referenced in the Prevailion post, which ushered us into seeing overlaps within the IOCs in the Positive Technologies blog. We found that these IOCs also overlap with those of the Azure-Sentinel detection rule for the APT41 threat actor group.

When we looked deeper into the activities of the threats within these clusters, the similarities continued. Reports from Subex and Positive Technologies described campaigns using PDF files that lured people in with a variety of tactics, including leveraging people's desire to see information indicating a swift end to the COVID-19 pandemic.

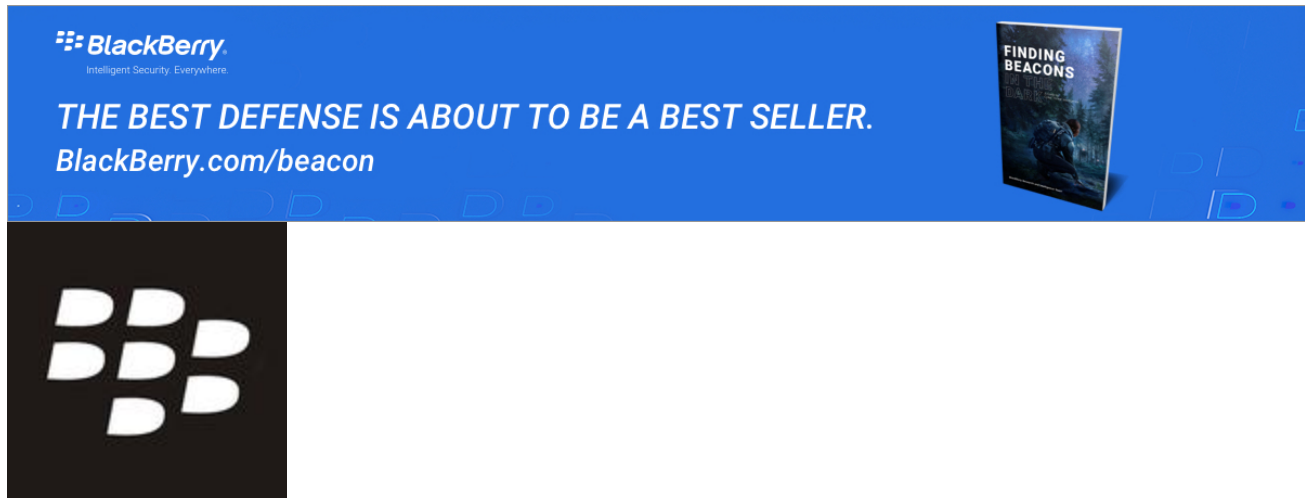
With the resources of a nation-state level threat group, it's possible to create a truly staggering level of diversity in a threat infrastructure. And while no one security group has that same level of funding, by pooling our collective brainpower we can still uncover the tracks that the cybercriminals involved worked so hard to hide.

Indicators of Compromise (IoCs)

Please view our GitHub for the IoCs referenced in this report: <https://github.com/blackberry/threat-research-and-intelligence/blob/main/APT41.csv>

Want to learn more about cyber threat hunting? Check out the BlackBerry Research & Intelligence Team's new book, *Finding Beacons in the Dark: A Guide to Cyber Threat Intelligence* - now available for pre-order

[here.](#)

A blue banner featuring the BlackBerry logo and tagline 'Intelligent Security. Everywhere.' on the left. The central text reads 'THE BEST DEFENSE IS ABOUT TO BE A BEST SELLER.' followed by the URL 'BlackBerry.com/beacon'. On the right, there is a 3D rendering of the book 'FINDING BEACONS' by the BlackBerry Research & Intelligence Team. Below the banner is a large black square containing the white BlackBerry logo.

About The BlackBerry Research & Intelligence Team

The BlackBerry Research & Intelligence team examines emerging and persistent threats, providing intelligence analysis for the benefit of defenders and the organizations they serve.

[Back](#)