infoblox

Threat Intel

MUDDLING MALSPAM: THE USE OF SPOOFED DOMAINS IN MALICIOUS SPAM

Authors:
Stelios Chatzistogias
Laura da Rocha
Renée Burton

TABLE OF CONTENT

This may read at first like a story about failed research, but it's actually a story about how research may seek one thing and wind up discovering something entirely different.

In March 2024, we published a blog report on an actor we call Muddling Meerkat, who conducts puzzling DNS operations via the Chinese Great Firewall. We had invested significant time in our research but were unable to figure out the purpose of these multiyear operations. Rather than putting the work into a drawer, we decided to release what we knew about the activity so others would share their own insights and, collectively, we might come to understand the true nature of Muddling Meerkat. It worked! The blog drew ideas from professionals in networking and security alike; some were able to provide anonymized data about their own view of Muddling Meerkat, or at least the so-called "target domains" we see in DNS.

Many of the suggestions for further research centered on spam operations. Some organizations had received abuse notifications for domains they owned, typically internal domains that were not used externally. The abuse reports were proof of large-scale spam distribution to big mail vendors like Google and Yahoo, and overwhelmingly, the source IP of the spam was assigned to China. This seemed consistent with Muddling Meerkat activities, in which we saw fake mail server (MX) records emanating from Chinese IP space, as well as similar MX queries coming into corporate networks through open resolvers.

One of the data files shared with us led to an epiphany: we owned several Muddling Meerkat "target" domains ourselves! That meant we could use abuse reports sent to us for these domains, as well as DNS authoritative name server logs, to better understand spam-related activity from a DNS perspective. But we also have a good spam collection ourselves, and we could hunt for campaigns that showed Muddling Meerkat behavior over time.

This paper is the result of our spam hunt. To be honest, we aren't sure if we are any closer to understanding Muddling Meerkat, which at first blush might be considered a failure. But in following those threads, we instead learned a lot about the use of domain spoofing in modern malicious spam (malspam) campaigns. We are going to share a few of our "catches" that show the more interesting ways that actors are employing domain spoofing today, all of which use some Muddling Meerkat-type behavior. We were able to connect these campaigns with the abuse reports we received from recipients and our authoritative DNS logs. Moreover, because we own some of the spoofed domains, we captured some of them as bounces back to our mail servers. By pivoting back and forth in these sources, we also learned more about the breadth of Muddling Meerkat target domains, expanding our original reported set from about 20 to over 650 domains.

Most surprising is just how pervasive domain spoofing is in spam. There are several mechanisms designed to protect users from spam in general and spoofing in particular, but we discovered that spoofing is still widely used. Most of the campaigns are sent from Chinese IP addresses, and the breadth of campaign types is quite remarkable. Despite security safeguards, the use of spoofed domains still pays off financially. In this paper, we'll look at:

- Modern campaigns that leverage QR codes in PDF attachments to steal from Chinese citizens,
- Popular brand impersonation targeting Japanese users to steal login credentials,
- Old extortion campaigns, possibly driven by botnet remnants, that attempt to con users into paying into the threat actor's crypto wallet, and
- Mysterious financial campaigns that appear to have no malicious content but also no motive.

In addition, we'll describe how we used our own authoritative DNS server logs to attempt to understand Muddling Meerkat but instead caught these spam campaigns.

DOMAIN SPOOFING IN SPAM

Threat actors can fake (spoof) the sender address of an email. They do this to make the email appear more legitimate. By using a domain that has been registered for many years, they are more likely to get past security mechanisms that check the sender domain age to identify malicious spam. On the other hand, if the actor spoofs a well-known domain such as amazon[.]com, there are several mechanisms the receiving mail server can use to determine when an email using one of these domains has been spoofed. We believe this risk of detection is why spammers are using old, neglected domains—the very same type of domain that Muddling Meerkat favors for their operations.

When a mail server receives email, it will perform several checks in DNS to attempt to validate the sender. It will then compare those results to the email headers. These checks include actions like verifying that the IP address from which the email was received is authorized to send email for that domain. Some of these checks rely on specific DNS records which often don't exist for old, neglected domains, and may result in a "soft" failure.

After the server carries out the standard checks and perhaps applies additional mail security algorithms, the email might be marked as spam or even quarantined. In other cases, it might make it through to the user's inbox. The malspam actor is hoping that their synthetic emails make it past enough spam traps to reach users and reap rewards.

AUTHORITATIVE DNS SERVERS AND SPAM

We happen to own some disused domains that have not actively hosted content for nearly 20 years. They lack most DNS records, including those that are typically used to check the authenticity of a sender domain, e.g., Sender Policy Framework (SPF) records. The domains are short and in highly reputable TLDs: perfect for Muddling Meerkat and spammers alike.

Ironically, several of our old domains are commonly cited in, for example, Tranco's top 1 million domain list. We suspect their popularity is driven entirely by spam. Without veering too far off the main topic of this blog, the popularity of our dormant domains is a great illustration of one of the reasons why top-list rankings need to be taken with a grain of salt. We have spent a lot of time studying domain popularity and threats; check out our previous papers.[1,2] (To read footnotes, please view this PDF online.)

DNS gives us a unique view of the abuse of our domains. We log queries for all our domains on our authoritative DNS server. These logs give us a window into a wide range of DNS activity, from internet scanning to spam distribution. In the case of email, a recipient's mail server will make several DNS queries to the authoritative server for the sender domain, including DNS TXT records. From our logs, we can see the IP address of DNS resolvers used by those mail servers and get a sense of the geographic distribution of the spam that is spoofing our domains.

We have also set up DomainKeys Identified Email (DKIM) records so that providers who receive spam from our domains can send us abuse reports via email. Those abuse reports include the IP address of the spam sender and timestamp information. We can combine them with the DNS TXT record requests to gain a pretty good view of how we are being falsely associated with spam distribution. Our mail servers don't transmit email, they only receive it.

Since we were interested in potential spam activity from Muddling Meerkat, we needed to isolate potential actor queries from others. There is a lot of noise in DNS. Many research organizations, like our own, make DNS queries to gather information and create a synthetic footprint in historical records. Our servers shouldn't receive any DNS queries because all the domains are dormant, yet they receive thousands of queries each day, sometimes tens of thousands. Figure 1 shows a comparison of the number of queries received at our authoritative server for four different Muddling Meerkat domains over time. The graph on top is for all record types, and the bottom one is for MX queries. These timeline charts indicate that mail-related activity is not necessarily correlated with overall DNS activity for the domains.

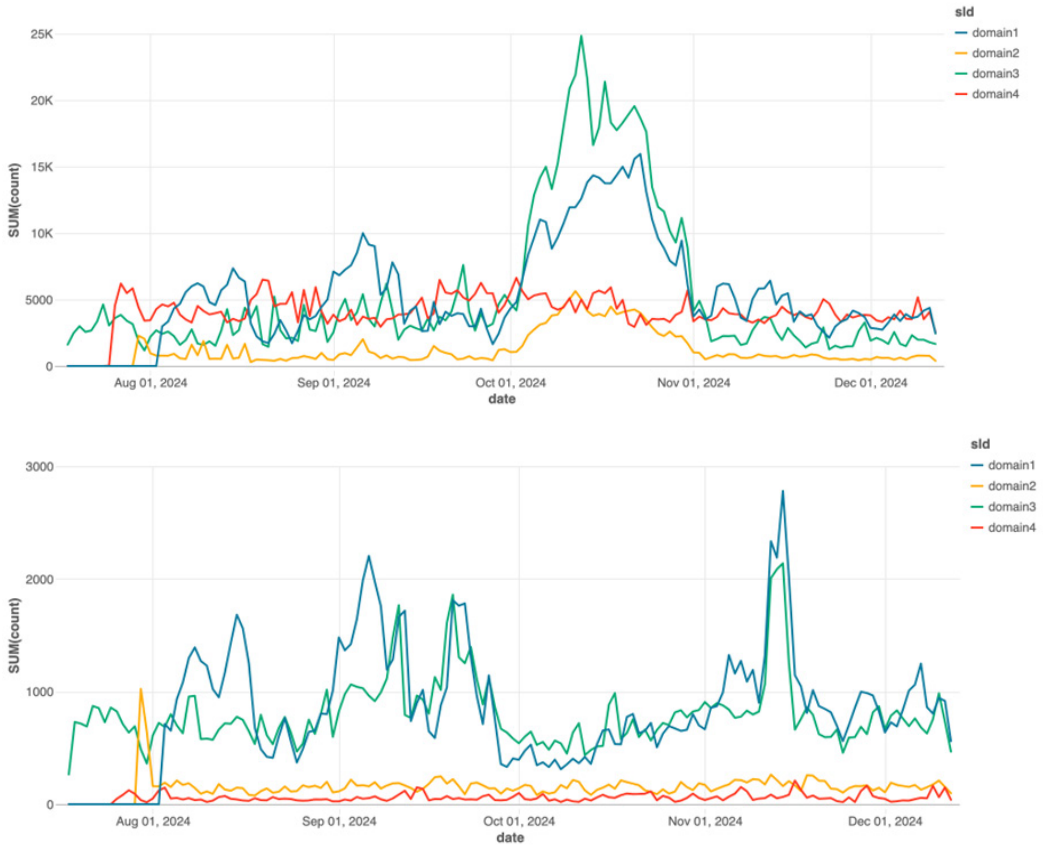

Figure 1. Top chart: query volume for all DNS record types at our authoritative server by domain; bottom chart: query volume for MX records

Most of the queries we receive at our authoritative servers do not match the Muddling Meerkat patterns, so we used various fingerprints built on prior research to isolate potential activity driven by the actor. We also compared those findings to the abuse reports we received via email. Muddling Meerkat DNS queries use different record types, but the most unusual from an investigative standpoint are MX record queries for short random subdomains. In the following example, if the target domain is target.domain, the query would look like:

```
<rand>.target.domain
```

The term "target" here is a loose one, as we explained in our earlier paper[3] because the actor targets these domains for use in their campaigns, rather than targeting them as part of an attack on the domain owners; the actor abuses these domains they don't own for an unknown purpose. We limited our analysis of queries to those that had hostnames that were only seen as a subdomain of a single domain we served and looked for trends. The length of uniquely observed hostnames varied, but those that were three characters long were most prevalent; see Figure 2. This was consistent with the data we had received from other domain holders. We also verified that queries came from large mail providers, such as Google, and mail security providers like Proofpoint.

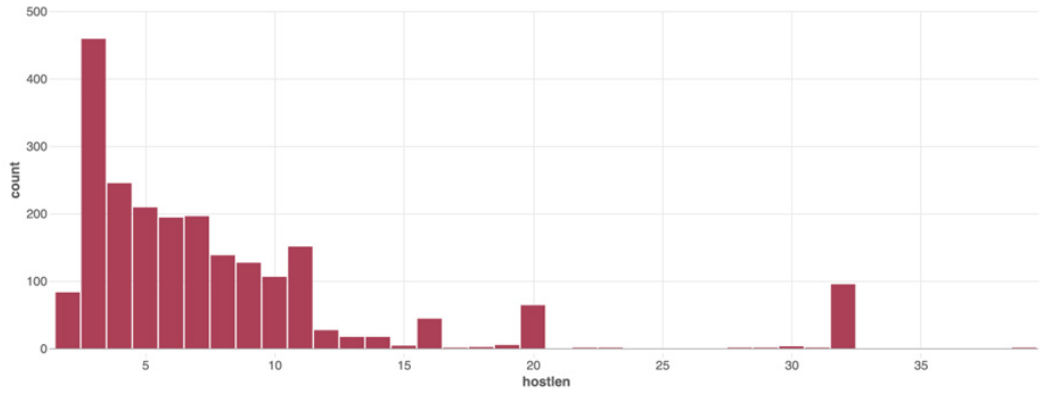

Figure 2. The length of uniquely observed hostnames in MX queries at our authoritative DNS servers

Empowered with the knowledge that our own domains were being used by Muddling Meerkat and were being spoofed by threat actors conducting malspam campaigns, we went hunting in our spam traps for active campaigns.

CATCH #1: QR CODE PHISHING CAMPAIGNS

The largest group of phishing campaigns that we observed spoofing our old domains targeted residents of greater China. These campaigns have run persistently since at least late 2022 and distribute attachments that contain a QR code that leads to a phishing site; see Figure 3. Based on our DNS data, abuse reports, and collateral information, we believe the attacks originate in greater China. The campaigns leverage a tactic that involves having the recipient open the email attachment and use WhatsApp to scan a QR code within. This two-step method creates additional challenges to securing users, because the attacker draws the victims from their laptops to an encrypted chat app, circumventing many common security measures. The threat actors also employ registered domain generation algorithms (RDGAs) to create random domains that are active for only a short period of time.

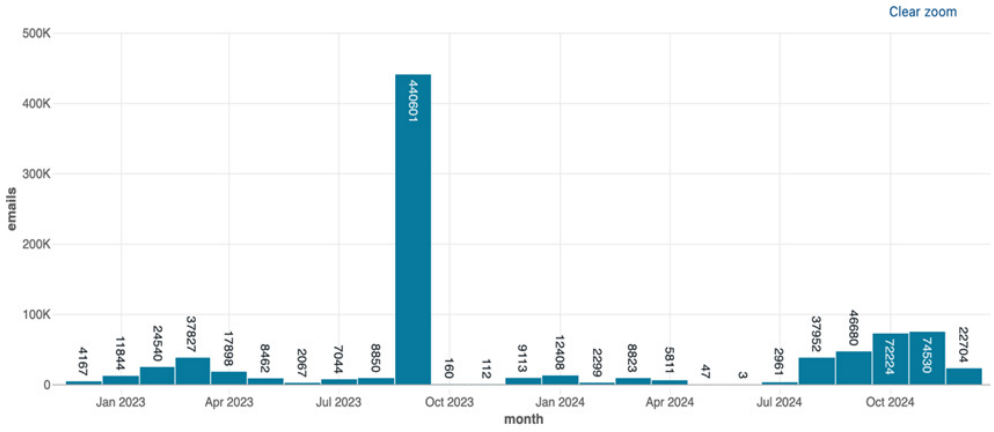

Figure 3. Volume of Chinese QR code phishing emails over time

These malspam campaigns use spoofed sender domains that include a large number of confirmed Muddling Meerkat target domains, including domains we own. Through spam analysis of these campaigns and by comparing historical DNS records, we extended the number of known Muddling Meerkat target domains from approximately 20 by March 2024 to over 650 today. The QR code campaigns, however, also contain many domains that Muddling Meerkat may well be using, but that we can't confirm through DNS.

These campaigns use sender email addresses that had a structure matching what we observed in Muddling Meerkat DNS queries. The sender's username was a short, random string of the form `<rand>@spoofed[.]domain`. Table 1 shows a sample of the campaign sender email addresses over time. Domains such as `jx[.]com` and `hm[.]com` were already known to be Muddling Meerkat target domains.

dm@**jx**[.]**com**	ino@**jjnywnd**[.]**com**
ab@**hm**[.]**com**	gwhy@**isathtooy**[.]**net**
zb@**iizlopn**[.]**com**	atmrp@**kym**[.]**net**
mu@**ibqg**[.]**net**	qivlzn@**kt**[.]**com**
xzu@**iejzhopjx**[.]**org**	atmrp@**kym**[.]**net**
iud@**irnvasa**[.]**net**	

Table 1: Sample of sender addresses for the QR code campaign; the sender emails have the pattern <2-9 random chars>@<spoofed[.]domain>

The email typically includes a tax-related lure in Mandarin and began in December 2022 or earlier. These appear to originate from Chinese IP space, primarily 4134 (Chinanet) and 56046 (China Mobile). Figure 4 shows some of the email subject lines and their English translations.

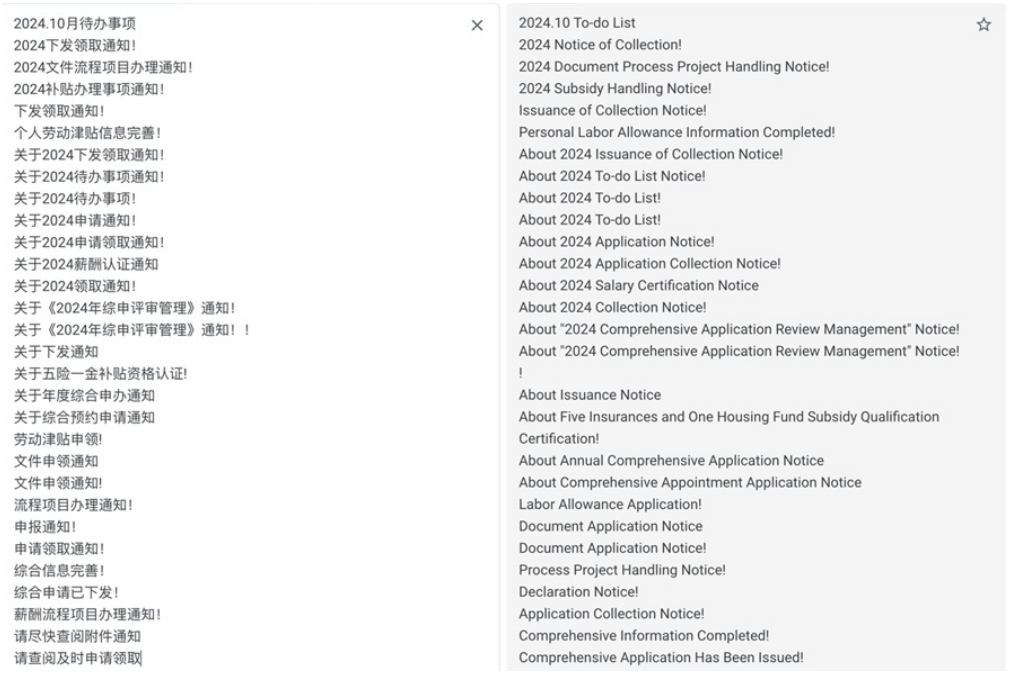

Figure 4: Sample of translated QR code campaign email subjects

Another distinguishing feature of this malspam is that most of the QR code documents are encrypted with a four-digit password, which is included somewhere in the email body, but not in any consistent way. Sometimes they are in parentheses, or they can be enclosed by other symbols. Figure 5 shows two examples of how passwords are included in the emails.

Figure 5. Two examples of how the four-digit password may vary in value and format in different emails; the red and green boxes highlight different ways the password is present in an email

The attachments contain a QR code with an embedded logo and instructions for the recipient to use AliPay/WeChat to scan the document; see Figure 6. These emails are no different than those we see cybercriminals use around the world to prey on vulnerable populations with promises of subsidies and financial benefits.

Figure 6. File attachment content and translation; the highlighted portion includes the instruction to use Alipay/WeChat to scan the QR code

Users on Twitter have reported the deception. According to the tweet[4] in Figure 7, one user was asked to enter a card number and identification information after scanning the QR code. They were then asked to enter the amount and the verification code which they assumed was for a payment to their account. Shortly after, a text message alerted them that they had paid 590 Euros from their card to the attacker. That is a significant return on a spam email!

Figure 7. Tweet from a user who was fooled by the QR code phishing scam

This scheme relies on second-stage phishing domains with a very short lifespan and that appear to be geofenced. They do not resolve in DNS after about a day and are in commonly abused TLDs like sbs, shop, life, bond and cn. These domains are made up of a random set of characters, e.g., `aaaefiuibew[.]cn` or `6ttox81[.]sbs`.

We are unable to say if this activity is from Muddling Meerkat. It seems more likely to be a common phishing-as-a-service (PhaaS) system. Although the campaigns do use the neglected domains we see with Muddling Meerkat, they appear to broadly spoof random domains, even ones that do not exist. The actor may use this technique to avoid repeated emails from the same sender. Despite efforts to protect users from malicious spam, some of these spoofs get through and are clearly profitable enough to maintain.

Other sender domains seen in this activity are shown in Table 2.

len2	len3	len4	len5
jt[.]net	iac[.]com	idhs[.]org	ivkpc[.]net
hc[.]com	izr[.]com	jxrn[.]org	jbdct[.]net
kk[.]net	koh[.]com	jirh[.]org	jfctl[.]org
jg[.]com	jwq[.]org	ismh[.]com	irnpc[.]net
kx[.]com	kcy[.]org	ikat[.]com	lahuf[.]net
len6	**len7**	**len8**	**len9**
jxjfwz[.]net	kbgpnek[.]org	jqmyuxk[.]com	hfababhqf[.]org
jxnsdf[.]net	ipcwfrn[.]com	jwruoytd[.]org	jfrcjfqjr[.]com
jwnlhr[.]org	iouwttz[.]com	ktfnmbxa[.]org	jkdduscaj[.]net
kindhy[.]net	jhrzbuk[.]org	jlsiwslr[.]org	jkjiwbpki[.]com
khznrl[.]com	hrggzxa[.]com	hrfliqoj[.]net	kwbjjlygw[.]net

Table 2. Sample spoofed domains seen in QR code phishing campaigns

Once we realized that the QR code campaigns spoofed domains that were outside of what we expected from Muddling Meerkat, we headed back to DNS and our spam collection to look for different campaigns that might be conducted by the Muddling Meerkat actor.

CATCH #2: JAPANESE PHISHING CAMPAIGNS

At our authoritative DNS servers, we noticed that an unusually large percentage of the mail-related queries included three-letter hostnames. As we tried to separate what queries might be created by Muddling Meerkat from those attributable to scanners and other sources, the volume and consistency of these queries felt like a good avenue of investigation. As a result, we looked for evidence of spam that had the same query structure.

We found a series of campaigns targeting Japanese users with emails that referred to popular brands such as Electronic Toll Collection (ETC, used on highways across Japan), Sumitomo Mitsui Banking Corporation (SMBC, one of the largest banks in Japan), as well as Amazon and Mastercard. The emails urge the user to authenticate with the service due to a security concern or other problem. A button included in the email leads the user into a traffic distribution system (TDS) and redirects them to a fake login page if certain criteria are met.[5] This method is common in malvertising and is used to cloak the final landing page to avoid detection by security companies. The fake login page steals the victim's credentials when they are entered. Figure 8 shows an example of these spam emails.

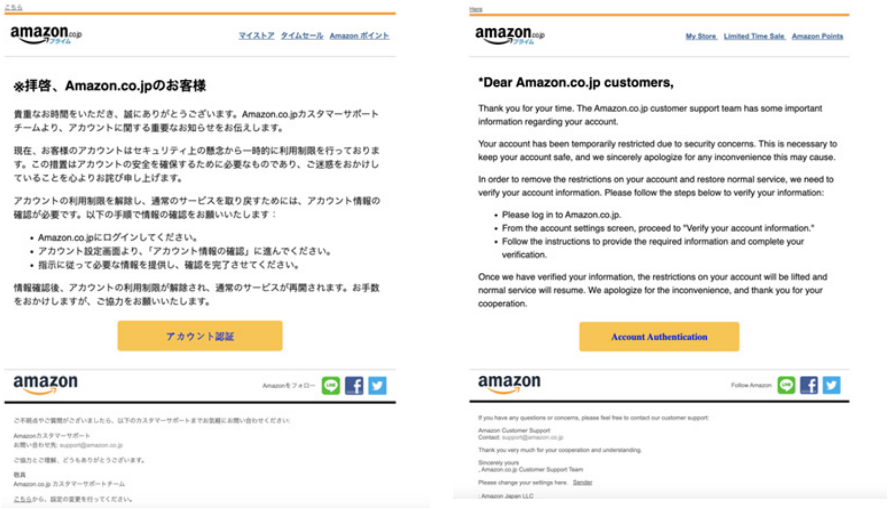

Figure 8. Original spam email example targeting Japanese users with fake Amazon warnings and a machine translation of the email

Once the user clicks on the "Account Authentication" button:

- They will be directed to `ubrjubf[.]com`, resolving on IP `43.128.150[.]42`
- Then the user will get redirected to a different domain `unpwple[.]com`, resolving on IP `43.133.182[.]243`
- The user arrives at a fake Amazon account login page; see Figure 9

Figure 9. Fake Amazon login page; image reference: https://urlscan.io/result/5c9bbf63-883f-4eab-b4fc-45e2809a8ac2/

We have observed several variations of Amazon-themed spam, as well as lures using Mastercard and SMBC Vpass.[6] This actor uses dedicated hosting infrastructure and rotates campaigns through the same domains and IP addresses.[7] The two dedicated IP addresses observed were `43.128.150[.]42` and `43.133.182[.]243`. Table 3 provides a list of RDGA domains used in the campaigns.

43.128.150[.]42	43.133.182[.]243
eujsubf[.]com, eujsxikw[.]com,	anzcinf[.]xyz, anzconc[.]xyz,
ikhcok[.]com, insjibr[.]com,	infkokf[.]com, omfkiht[.]xyz,
insjkf[.]com, khcpw[.]com,	omfkybg[.]xyz, inybinf[.]com,
maczplw[.]com, maczunf[.]com,	unpwple[.]com, inybubf[.]com,
pknribt[.]com, pknrinf[.]com,	pplaaej[.]com, eccteukx[.]com,
pknrinr[.]com, pknrohv[.]com,	espoeubf[.]com, unpwmlw[.]com,
pknrybg[.]com, pknrynf[.]com,	pplaaeu[.]com, ecctenje[.]com,
ubrjpnf[.]com, ubrjubf[.]com,	unpwibr[.]com, ecctepje[.]com,
unpwinf[.]com, uwkxubs[.]com,	pplaaep[.]com, pplaaea[.]com,
wkxaubf[.]com, wkxaunf[.]com	espoeunf[.]com, espoekwl[.]com

Table 3. Sample of RDGA domains on dedicated IP addresses used in campaigns targeting Japanese users

Like the campaigns described in the previous section, the emails in these campaigns use spoofed sender domains, including domains owned by Infoblox Threat Intel. They also follow the format we found prevalent at our authoritative DNS servers with a three-character subdomain, and in the abuse reports we received from mail providers. Table 4 shows a sample of sender email addresses.

ak@fdd.**xpv[.]org**	iipnf@gvy.**zxdvrdbtb[.]com**
mh@thq.**cyxfyxrv[.]com**	zmrbcj@bce.**xnity[.]net**
mfhez@shp.**bzmb[.]com**	nxohlq@vzy.**dpyj[.]com**
gcini@vjw.**mosf[.]com**	

Table 4. A sample of sender addresses for Japanese phishing emails with three-letter subdomains; the three-letter hostnames are colored in red, while the spoofed domain is in bold

This wasn't the only type of campaign we saw targeting Japanese users. Another major lure included MyEtherWallet, a popular crypto wallet, and used lookalike domains. The spam messages sometimes include Japanese text, e.g., "(重要なお知らせ) MyEtherWallet ご利用確認のお願い", which translates to "[Important Notice] Request for confirmation of use of MyEtherWallet," and ask users to login to their account. See Figure 10 for an example of an English-language email. Although the link appears to be the real website, it actually leads to a lookalike domain created by the threat actor.

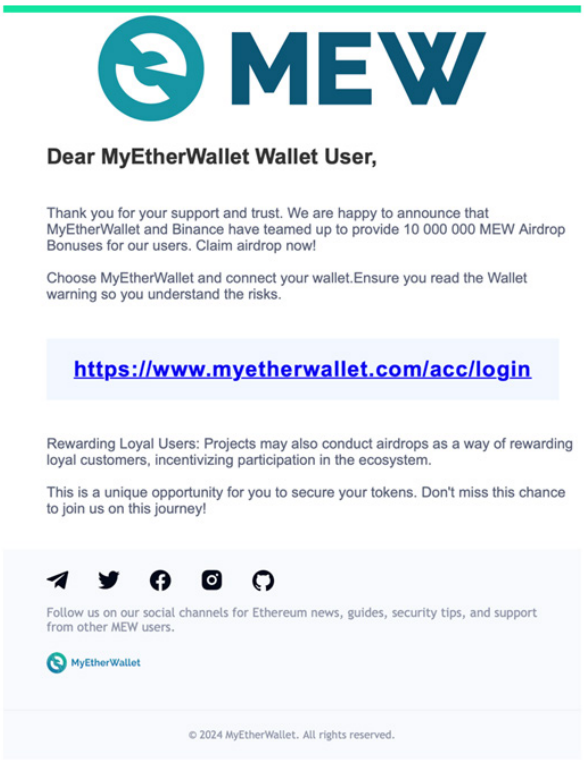

Figure 10. Sample spam campaign targeting Japanese users; this particular email had the subject line "Binance Distribution of MyEtherWallet (MEW) Airdrop" and led the user to `myetherwallatak[.]org`

The lookalike domains led to a mirror copy of the MyEtherWallet website and were used to steal user credentials. These domains are in several different TLDs including com and org; see Table 5 for a sample.

myetherwalletie[.]com	myetherwalletih[.]com	myetherwalletik[.]com
myetherwalletiv[.]com	myetherwalletjp[.]com	myetherwalletrt[.]com
myetherwallettv[.]com	myetherewallet[.]org	myetherswallet[.]org
myetherwallata[.]org		

Table 5. Sample lookalike domains used to phish credentials from Japanese users

In the QR code phishing campaign, the Muddling Meerkat domains were seen in the so-called sender addresses, i.e., the email addresses visible to the recipient. However, in this Japanese variant, the domains are seen in the "received from" portion of the message, which is used for the technical delivery via SMTP. We found ourselves with another set of malicious spam campaigns that use the same domains that are observed in Muddling Meerkat operations, including a similar subdomain format, but we are not able to verify that they are the work of Muddling Meerkat. Table 6 provides a sample of the spoofed domains.

len2	len3	len4	len5
xl[.]com	tgt[.]org	iddm[.]org	fhqqc[.]com
gz[.]net	paf[.]org	yqqb[.]org	mseur[.]com
ed[.]org	bcc[.]com	nvso[.]net	ofddy[.]com
df[.]org	zla[.]com	nqui[.]com	agejx[.]net
wx[.]com	tgf[.]net	duth[.]net	mlngi[.]com
len6	**len7**	**len8**	**len9**
kwwezx[.]net	gbiutoj[.]com	mitsxpjh[.]com	nxbfvjkhr[.]org
bwidqv[.]com	jeihdgt[.]com	wtfmbcvt[.]com	lkhyleslk[.]net
piuxic[.]com	qspdwhc[.]com	jgggzbmq[.]org	nmshofzmh[.]net
xdgzas[.]com	grjfgpw[.]net	qqegowhv[.]org	ykbhnoers[.]com
nwfffu[.]org	vudgfcm[.]net	invphyzf[.]com	mqnbqsygn[.]org

Table 6. Sample spoofed domains seen in MyEtherWallet campaigns

One of several gaps in understanding these campaigns is that the number of domains we've discovered is likely too small of a sample set to validate that all the domains overlap with Muddling Meerkat. However, this set of campaigns is another example of how China-linked threat actors are using domain spoofing for their spam operations.

With that, we headed back to the spam traps.

CATCH #3: FAMILIAR EXTORTION CAMPAIGNS

Not only did we find domain spoofing in the QR code and Japanese-language campaigns, but we also found it in campaigns that leveraged well-known spam tropes. Extortion emails claiming a hacker has accessed the user's device and recorded some embarrassing activities are a staple in the malspam world. We were a bit surprised to find that these also use spoofed sender domains, but with a twist: the actor spoofs the user's own email address and challenges them to check it and see. The email tells the user that their device has been compromised, and as proof, the actor alleges that the message was sent from the user's own account. And yet they hadn't; the mail headers show it came through Chinese IP addresses, not the user's. See Figure 11 below for an example of the email content.

The email tells the user to pay the sender in return for removing the malware from their device and includes a Bitcoin wallet address, which varies across spam messages. We don't know if this is an extortion service or if the same actor is using a variety of wallets. The victims are asked to pay US$1800 in the examples we have collected. While it may seem surprising that many users would actually read these spam emails, much less act on them, the scam apparently works. Checking the balance of these wallets with `bitref[.]com`, we can see that they do contain significant funds; one wallet contained nearly US$26k.

> Hello there! Unfortunately, there are some bad news for you. Some time ago your device was infected with my private trojan, R.A.T (Remote Administration Tool), if you want to find out more about it simply use Google. My trojan allowed me to access your files, accounts and your cam. Check the sender of this email, I have sent it from your email account. To make sure you read this email, you will receive it multiple times. You truly enjoy checking out porn websites and watching dirty videos, while having a lot of kinky fun. I RECORDED YOU (through the cam of your device) SATISFYING YOURSELF! After that I removed my malware to not leave any traces. If you still doubt my serious intentions, it only takes couple mouse clicks to share the video of you with your friends, relatives, all email contacts, on social networks and the darknet. All you need is $1800 USD in Bitcoin (BTC) transfer to my account. After the transaction is successful, I will proceed to delete everything. Be sure, I keep my promises. You can easily buy Bitcoin (BTC) here: https://cex.io/buy-bitcoins https://nexo.com/buy-crypto/bitcoin-btc https://bitpay.com/buy-bitcoin/?crypto=BTC https://paybis.com/ https://invity.io/buy-crypto Or simply google other exchanger. After that send the Bitcoin (BTC) directly to my wallet, or install the free software: Atomicwallet, or: Exodus wallet, then receive and send to mine. My Bitcoin (BTC) address is: 1GtGZpzfRkAVBL48F68mi8bTcatwpTZGm8 Yes, that's how the address looks like, copy and paste my address, it's (cAsE-sEnSEtiVE). You are given not more than 3 days after you have opened this email. As I got access to this email account, I will know if this email has already been read. Everything will be carried out based on fairness. An advice from me, regularly change all your passwords to your accounts and update your device with newest security patches.

Figure 11. An example of extortion spam that leverages spoofed sender domains

It seems likely these campaigns, and possibly many others using spoofed sender domains, are originating from lingering spam bots. Minimally, the attackers aren't validating the victims' email addresses to ensure they are received or read. We have instances where the recipient's email address was tied to one of our domains that last hosted content in 2007 and had not had email users in over 15 years. There are no breach records that might explain why these emails were triggered, and it's unknown to us whether, in fact, these users ever existed.

This, and other similar spam campaigns we found, conjure up images of abandoned spam cannons left to drift in internet space. We also saw old worms being transmitted, another sign of botnet remnants left to run while malicious spammers moved on to techniques like the QR codes and fake account pages like those we've shown above. These campaigns, now likely on autopilot, seem to be more likely echos than the more recent work of a sophisticated actor like Muddling Meerkat.

CATCH #4: MYSTERIOUS MALSPAM

This whole research agenda began with a mystery, and we'll end this paper with another: a very active spam campaign that uses spoofed sender domains and includes seemingly innocuous Excel spreadsheet attachments that have no evident purpose. We can't explain the motive for these emails, which spoof the same types of domains that Muddling Meerkat uses.

These emails purportedly come from **上海亚凯**, which translates to "Shanghai Yakai," the name of a Chinese freight company. The email addresses differ widely and include synthetic usernames like "Edward.Evelyn" and "Heidi.Gracie." Campaigns were seen every two out of three days in 2024, but didn't vary. The subject line indicates that the email contains new freight rate updates, and the attachment is a single named spreadsheet: **上海亚凯国际运价表**.xlsx. We have found no malicious content in these files.

There is no Call to Action (CTA) in the email. By all appearances, it is just a continually updated set of freight rates for a Chinese shipping company. But for what purpose? These emails do not appear to be sent to customers who have forgotten to change their email address or unsubscribe. The use of domain spoofing removes any sense of legitimacy, and it seems unclear why either a shipping company or a malicious actor would send emails like these. Table 7 shows a sample of the sender domains.

len4	len5	len6	len7
igeb[.]net	accou[.]com	axegal[.]com	awpking[.]com
kwfm[.]com	drsmj[.]com	devsmx[.]com	comitis[.]com
pqhh[.]com	eddim[.]com	glypix[.]com	donmenn[.]com
rrbc[.]com	hetoo[.]com	gulart[.]net	fundsle[.]com
tkee[.]net	horek[.]com	jomila[.]net	karnege[.]com
tnmc[.]com	memsz[.]com	mzylla[.]com	mtrplay[.]com
ukei[.]net	svard[.]net	okayme[.]com	rajprem[.]com
utpz[.]com	tapli[.]net	theiwl[.]com	techsox[.]com
vbhh[.]com	uweko[.]com	vaites[.]com	tjipbpo[.]com
wuwo[.]com	youbi[.]com	ynglet[.]com	wulthur[.]net

Table 7. A sample of spoofed sender domains used in Shanghai Yakai freight spam

A similar campaign technique was seen in personal spam, but instead of messages from a freight company, the email provides mutual fund values from an Indian investment company. These messages, which are flagged by Google Mail as suspicious spam, also contain an innocuous spreadsheet and a PDF file. In this case, the username of the sender is a former acquaintance, and it seems likely their email account was hacked at some point for use in spam operations. But like the Chinese freight spam, it is unclear how these messages have value for the spam actor.

VIEW FROM THE AUTHORITATIVE DNS SERVER

Muddling Meerkat has conducted strange DNS operations for over six years. These involve fake responses from the Chinese Great Firewall and the use of long-neglected domains that they don't control. While their DNS activity includes several record types, the fake responses are for MX records of the base, or target, domain. For example, DNS responses containing MX records for `kb[.]com` are observed from Chinese IP addresses, even though `kb[.]com` has no MX records. Moreover, these fake records include a short, random hostname that is only observed once over time—e.g., `x4rd.kb[.]com`—which might be an observed MX record for `kb[.]com`. When we first published in March 2024, we had identified about 20 such domains but now have confirmed several hundred others.

In addition to looking for evidence of spam operations from the actor, we also analyzed DNS logs at our authoritative servers and attempted to match them to the fake DNS responses observed in collateral data for the domains we owned. The hypothesis was that if we could see a query for one of the fake MX record domains, e.g., `x4rd[.]our[.]domain`, we could use the querier's IP address to better understand Muddling Meerkat operations. Unfortunately, we were unable to definitively match the Muddling Meerkat records to queries at our servers.

What does it mean that this match cannot be found? Well, it implies that whoever or whatever receives the fake MX responses, e.g., `x4rd[.]our[.]domain`, does not use those responses in any follow-on DNS queries and does not appear to use them for spam. This lack of clear motivation seems to destroy the notion of a botnet receiving domains to use in spoofed emails. So, then, what are the responses used for? No idea. Muddling Meerkat remains a mystery. Got ideas or a different perspective? We're all ears.

CONCLUSION

We weren't able to determine what Muddling Meerkat is up to, but our investigation was ultimately successful: we learned a great deal about how actors use spoofed domains in malspam, which can inform ways to stop them. For threat researchers like us, that insight is often every bit as important as knowing the intentions behind them.

You can't always get what you want, but you just might find, you get what you need.[8]

INFOBLOX THREAT INTEL

Infoblox Threat Intel is the leading creator of original DNS threat intelligence, distinguishing itself in a sea of aggregators. What sets us apart? Two things: mad DNS skills and unparalleled visibility. DNS is notoriously tricky to interpret and hunt from, but our deep understanding and unique access give us a high-powered scope to zero in on cyber threats. We're proactive, not just defensive, using our insights to disrupt cybercrime where it begins. We also believe in sharing knowledge to support the broader security community by publishing detailed research and releasing indicators on GitHub. In addition, our intel is seamlessly integrated into our Infoblox DNS Detection and Response solutions, so customers automatically get its benefits, along with ridiculously low false positive rates.

Corporate Headquarters
2390 Mission College Blvd, Ste. 501
Santa Clara, CA 95054

+1.408.986.4000
www.infoblox.com

Version: 20241223v1