CYBER
THREAT
ANALYSIS

Recorded Future®

By Insikt Group®

January 18, 2022

2021 Adversary Infrastructure Report

Recorded Future's Insikt Group® conducted a study of malicious command and control (C2) infrastructure identified using proactive scanning and collection methods throughout 2021. All data was sourced from the Recorded Future® Platform and is current as of December 10, 2021.

Executive Summary

Recorded Future tracks the creation and modification of new malicious infrastructure for a multitude of post-exploitation toolkits, custom malware, and open-source remote access trojans (RATs). Since 2017, Insikt Group has created detections for 80 families, including RATs, advanced persistent threat (APT) malware, botnet families, and other commodity tools. Recorded Future observed over 10,000 unique command and control (C2) servers during 2021 across more than 80 families. Our collection in 2021 was dominated by Cobalt Strike Team Servers and botnet families, both of which applied more resiliency and stealth measures throughout the year.

Key Findings

- Our prediction last year anticipating an increase in Sliver, Mythic, Covenant, and Octopus C2 frameworks was only partially correct. While there has been small increase in use of Covenant, Sliver and Mythic, our visibility has shown continued reliance on Cobalt Strike with minimal adoption of newer C2 frameworks.
- 25% of detected servers (3,400 servers) were not referenced in open sources; they were only identified on the Recorded Future Command and Control source.
- Recorded Future observed an average of a 35-day lead time between when a C2 server is detected by our scanning efforts and when it is reported in other sources.
- While Emotet's return has garnered headlines, other botnets have continued to insulate, diversify, and grow their infrastructure during Emotet's absence in 2021.

Background

Lead time in identifying malicious servers can be a proactive measure in neutralizing threats. Before a server can be used by a threat actor, it has to be acquired, either via compromise or legitimate purchase. Then, the software must be installed, configurations tuned, secure sockets layer (SSL) certificates registered, and files added to the server. The actors must access it via panel login, secure shell (SSH), or remote desktop protocol (RDP), and then expose the malware controller on a port to allow the data to transfer from the victim and to administer commands to infections. Only then can the server be used maliciously.

However, in exposing, configuring, and accessing the server, the adversary leaves behind their fingerprints, which can be in software versions deployed on the server, the login panel, SSL registration patterns, or the default message returned by a simple probe. These fingerprints create detection opportunities before a phishing email is sent or an implant is compiled.

Similarly, such a collection can illuminate many things about adversaries. Seeing how many command and control (C2) servers are created can help quantify the breadth of malicious campaigns. Comparing such data to reports of intrusions related to those families can identify how many intrusions get caught and potentially how many events remain unknown in the public domain. Measuring the tempo of server creation can provide insight into forthcoming surges or drawdowns of activity. Finally, collection can provide novel indicators and intelligence otherwise unavailable in the public domain.

A Note on Collection Bias

Recorded Future collects information about C2 servers based predominantly on traits from known malware families and their server-side software. The nature of this collection through passive and active internet scan data will be focused on collecting known command and control frameworks and their derivatives or support infrastructure. Recorded Future cannot verify an IP address to be a C2 server without proof of malicious activity of at least one of the servers in a given family. Therefore, we will be biased in reporting servers of known threats and have a collection bias towards those servers. These should not act as a replacement for identifying anomalies or detecting odd traffic inside a network.

Threat Analysis

The most commonly observed families were mixed between post-exploitation frameworks and botnet infrastructure. Cobalt Strike Team Servers were again the most detected C2 controller, representing 23.7% of the total C2 servers identified. Similar to last year, Metasploit and Meterpreter represented the other top C2 servers identified by Recorded Future. TrickBot and QakBot were also among the 5 most detected families.

Top 5 Most Detected C2 Families	
Family	2021 C2s
Cobalt Strike Team Servers	3691
Meterpreter	396
Metasploit	710
QakBot	571
TrickBot	468

Table 1: Top detected C2 infrastructure by total unique servers

Post-Exploitation Frameworks

It is difficult to estimate what percentage of any post-exploitation frameworks we detected are used in legitimate red teaming operations, and which are used by criminal or espionage elements. Cobalt Strike has a diverse set of users, gaining the most direct attention in its use in ransomware operations. There was increased adoption of Mythic, Covenant, and high-profile use of Sliver during 2021; most other post-exploitation tools saw similar volumes of deployment compared to 2020. Increased volume may be tied to improved signatures and increased collection efforts, but is at least partly a reflection of continued adoption by various operations.

Detections of unaltered Cobalt Strike Team Servers (via the pre-configured TLS certificate, the presence of a Beacon payload, or telltale HTTP headers) represented 23.7% of the total C2 servers identified. This detection includes a small subset of C2s using domain-fronting, but most detected Cobalt Strike Team Servers are the base model; we do not have an estimate for the number of well-insulated or obfuscated Cobalt Strike Team Servers in use. The total number of detected Cobalt Strike Team Servers includes those with Malleable Profiles.

BlackBerry researchers observed roughly 6,000 unique Team Servers in 2021, based both on Beacon payload analysis and scanning operations. The gap in detections between pure scanning operations and a multi-faceted approach highlight how many Team Servers are evading identification (nearly 3,000) and how vital diverse collection and analysis functions are for defense.

In 2021, Microsoft, RiskIQ, and Insikt Group all identified instances of initial access brokers setting up ready-to-use Cobalt Strike Team Servers and infections to their clientele. This phenomenon was dubbed Cobalt Strike C2-as-a-Service. This does not include the various Cobalt Strike Team Servers used by various loaders or by ransomware affiliates to kick-start operations from other accesses.

The increased detection of Cobalt Strike is partially attributable to improved detection methods and analytics, as evidenced by the public proliferation of more aggressive collection mechanisms. Researchers continue to publish detection logic and findings that allow for the decryption of traffic from a large subset of Cobalt Strike Beacons.

Top 10 Observed Offensive Security Tools			
Family	2021 C2s	2020 C2s	Previous Notable Users
Cobalt Strike	3691	1441	APT41, Mustang Panda, Ocean Lotus, FIN7
Meterpreter	731	259	COBALT ILLUSION
Metasploit	710	1122	JointWorm (EVILNUM), Turla
Powershell Empire	269	289	Sandworm, GADOLINIUM
Covenant	180	51	GreenBug, FIN12
PupyRAT	177	454	MuddyWater, TA505
Sliver	169	27	WellMess Operators, TA551
Mythic	163	28	N/A
Koadic	109	19	Sofacy
PoshC2	103	12	UNC1945

Table 2: Most common offensive security tools by C2 servers detected in 2021

Going Back to Bots

The briefly successful law enforcement takedown of the Emotet botnet left a void in the loader and botnet market. Despite frequent breaks, the botnet was one of the most prolific and profitable threats in 2019 and 2020. Although other operations, such as TR distributor (ChaseLdr, TA577) or TA551 (Shathak), attempted to fill the spam gap left by Emotet's months-long disappearance, TrickBot, QakBot, Bazar, IcedID, and Dridex all continued onward. All of them have been observed acting as precursors to ransomware; IcedID was linked to Egregor deployments, TrickBot and Bazar families have been linked to Ryuk and Conti use, Dridex has led to DoppelPaymer, and QakBot has deployed ProLock and DoppelPaymer.

Sample over sample, the botnets varied in how many embedded C2 servers they referenced. TrickBot averaged 20 C2 servers per sample configuration, while IcedID and Dridex each averaged 3 servers, and QakBot averaged a much larger 142 C2 IP addresses per configuration. Resurfaced versions of Emotet averaged 20 servers per configuration.

Top 5 Most Prolific Botnet Families	
Family	**2021 C2s**
TrickBot	571
QakBot	516
Bazar/Baza Family	405
Dridex	383
IcedID	332

Table 3: Most common botnet families by C2 servers detected in 2021

The botnets vary greatly in terms of active operational size; sampling concurrently active servers from Recorded Future data showed TrickBot maintained 132 C2 servers at a given time, Bazar peaked at 99, while IcedID had 83, QakBot had 72, and Dridex had 32 active C2 servers. This measure does not account for how the operators divide the infrastructure but is based on observable operational size indexed within the Recorded Future Platform.

QakBot (QBot)

QakBot contains the largest list of C2s per sample, as shown by data from the configuration file included in the malware's resource section. QakBot's use of higher TCP ports partially hindered proactive scanning; roughly 3,200 total IP addresses were observed through configuration extraction, where 516 servers were observed via proactive scanning. Infrastructure throughout 2021 has also been divided across affiliates; for example, QakBot samples deployed by the TR distributor (ChaseLdr, TA577) and the Obama affiliate shared 95% of the same infrastructure, distinct from the 95% infrastructure overlap between the Biden and Clinton affiliates. The Abc affiliate operated on a relative island, with only 8% overlap with the Biden and Clinton infrastructure. Previously, TR and Biden affiliates had a large overlap in infrastructure; this cycling is likely to improve resiliency across the operation.

TrickBot

TrickBot's large volume of infrastructure for 2021 is partially accounted for by the botnet splitting its infrastructure into 2 distinct branches. Each branch was built by distinct server types in early 2021, before using identical but distinct C2 nodes beginning in March 2021. Recorded Future's visibility indicates the "original branch" serviced TrickBot clients while the "secondary branch" was used by TrickBot operators or highly trusted affiliates. The shift is likely a resiliency effort: TrickBot operators can continue spamming operations via a secondary botnet branch, even if the more used one gets taken down or blocked based on the volume of infections.

Bazar Family

Both BazarLoader (YerLoader, BazaLoader, and KEGTAP) and BazarBackdoor (Syndet, BEERBOT) often share infrastructure or deploy similar server configurations, leading us to list them together. Although the family can and originally did use Emercoin DNS (EmerDNS) .bazar as primary/backup domains for command and control, operators have often relied on more straightforward IPv4 addresses during 2021. The Bazar families, more so than others listed here, continually evolved their server software, deploying multiple server configurations simultaneously during 2021, likely in attempts to evade detection.

Dridex

In comparison, Dridex has a much smaller list of hardcoded C2s in its respective configurations. Each Dridex affiliate uses a different configuration file embedded in Dridex's main payload. The affiliate IDs are regularly rotated or aged off, making it difficult to track customers or methods used for distribution over time. Dridex has regularly updated its payload to make it more stealthy, but there are no indications that it did the same with its infrastructure in the wake of Emotet's takedown.

IcedID (BokBot)

Unlike other botnets, IcedID does not openly track its clients or users, although we believe its developers operate under an affiliate model. IcedID infrastructure is also less mature in its operational security, continuing to use the same pool of servers for its operations despite the publication of methods to ID the servers. Although IcedID effectively generates infections, it does not display the same level of operational security as the other families mentioned.

Emotet

Emotet's return has rightly caused concern among security researchers. Within 24 hours of new Emotet samples being loaded by TrickBot, the botnet began spamming. Within a week, a second branch was established to service more global spam operations. There have been 40 positively identified Emotet C2s derived from the samples across both epochs, with an additional 45 servers that share server patterns that will likely be used as Tier 1 servers by Emotet. There have been at least 4 servers observed hosting both Dridex and Emotet Tier 1 controller software. While the total count is not greater than other botnets mentioned here, the infrastructure creation rate indicates the operators have ambitions to return Emotet to its former prominence and power.

Global Scale

We observed the creation of C2 infrastructure on 1,650 hosting providers across 130 different countries. While this represents a majority of global geography, the abused servers account for only a small percentage of the total AS operators, which exceeds 60,000 providers. The data indicates the largest hosting providers are the most abused for C2 hosting; 20 AS operators (12% of total ASNs observed) had more than 100 C2 servers detected on them during 2021.

- The US hosted 4,654 C2 servers in 2021; China was second with 1,949, and Germany was third with 629.
- 858 AS operators (52% of total ASNs observed) were observed hosting 1 C2 server during 2021.
- 24 countries of the 130 observed hosting C2 servers hosted only C2 1 server during 2021.
- While Recorded Future observed 1,630 unique AS operators hosting C2 servers during 2021, the heavy majority were observed hosting 100 or fewer C2s.
- 1,454 AS providers (88% of total ASNs observed) hosted fewer than 10 C2 servers.

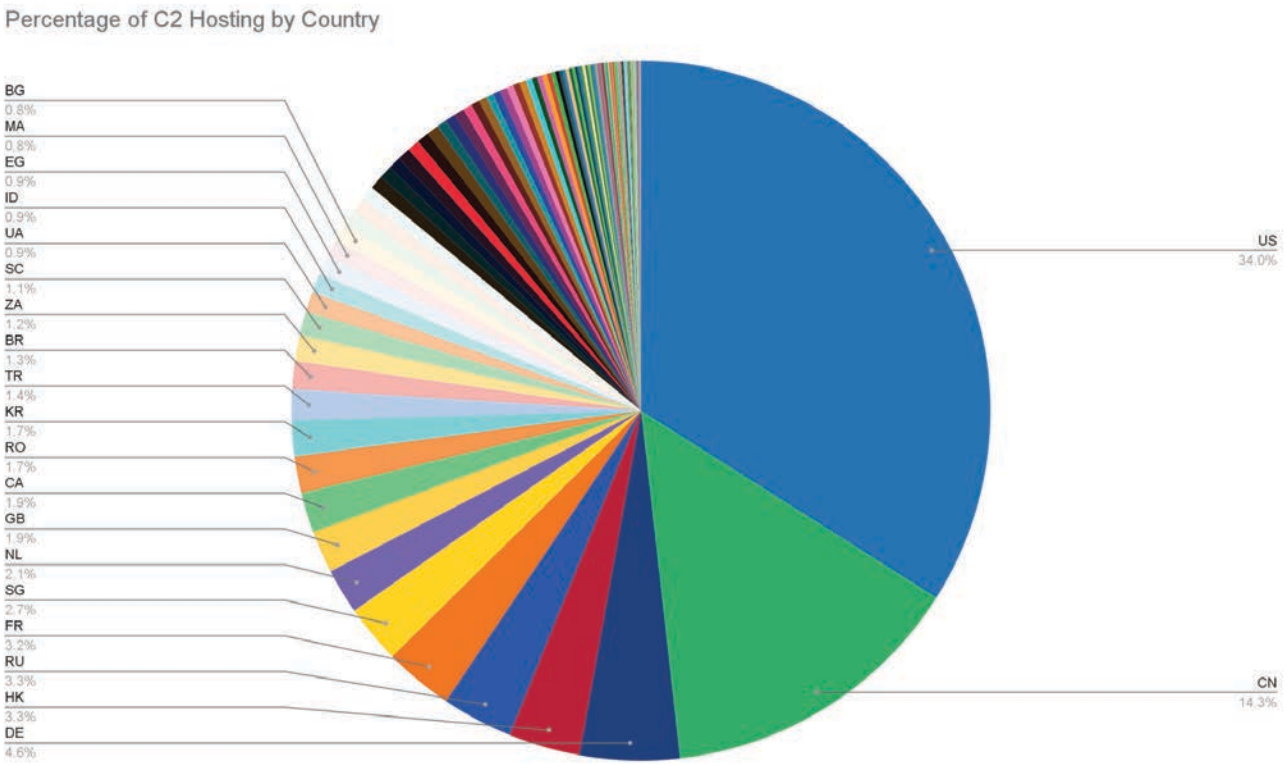

Figure 1: Most common countries based on total C2 servers detected in 2021 (Source: Recorded Future)

Top 10 C2 Hosting Providers By Volume				
Hosting Provider	**ASN**	**Country**	**Top Family**	**Total C2's**
Digital Ocean	AS14061	United States	Cobalt Strike	968
Choopa, LLC	AS20473	United States	Cobalt Strike	700
Amazon.com, Inc.	AS16509	United States	Cobalt Strike	624
Hangzhou Alibaba Advertising	AS37963	China	Cobalt Strike	574
Shenzhen Tencent Computer Systems	AS45090	China	Cobalt Strike	571
OVH SAS	AS16276	France	Cobalt Strike	267
Linode LLC	AS63949	United States	Cobalt Strike	208
Microsoft Corporation	AS8075	United States	Cobalt Strike	205
BGPNET Global	AS64050	Singapore	Metasploit	181
M247 Ltd	AS9009	United Kingdom	Cobalt Strike & PupyRAT	171

Table 4: Hosting providers who hosted the most C2 servers during 2021

Digital Ocean, operating out of the United States, hosted the most C2s of any of the ASNs observed by Recorded Future. They accounted for 968 individual C2 servers (7.1%). The most commonly observed family on Digital Ocean was Cobalt Strike, with 167 servers identified. The next largest was Choopa LLC (recently renamed the Constant Company), a Virtual Private Server provider operating out of the US, while the owner of the most C2s observed in 2020 (Amazon.com Inc.) dropped to third in 2021.

While these hosting providers accounted for the largest number of C2 servers, the C2 servers represented a minuscule percentage of total number of servers under their jurisdiction. The table below highlights the 10 providers with the highest percentage of C2 servers compared to their total holdings. This estimate is based on the number of IPv4 prefixes announced by the AS, compared to confirmed C2 servers observed in 2021.

Some of these providers, such as Media Land LLC, can be categorized as bulletproof hosting providers and are being marketed on underground forums. These providers also had much more diverse malware hosting than the largest volumes seen above, including AXIOMATICASYMPTOTE, a server-side component typically used to administer ShadowPad infections.

Hosting Providers with Highest Percentage of C2s Hosted				
Hosting Provider	**ASN**	**Country**	**Top Detection**	**Percent of Hosts are C2**
Media Land LLC	AS206728	Russia	Cobalt Strike	5.96%
Lider Telecomunicaçoes Eireli	AS268773	Brazil	njRAT	5.07%
Danilenko, Artyom	AS208476	Germany	AsyncRAT	2.60%
International Hosting Solutions LLP	AS213354	United Kingdom	BazarLoader	2.53%
NXTSERVERS SRL	AS64398	Romania	IcedID	1.41%
HOSTKEY	AS395839	United States	Cobalt Strike	1.07%
Beijing 3389 Network Technology	AS136146	China	Fakenocam	0.90%
RM Engineering LLC	AS49877	Russia	Cobalt Strike	0.72%
HDTIDC LIMITED	AS136038	Hong Kong	AXIOMATICASYMPTOTE	0.56%
Host Sailor Ltd.	AS60117	United Arab Emirates	Cobalt Strike	0.46%

Table 5: Hosting providers who hosted the highest percentage of C2 servers compared to total servers during 2021

Outlook

We anticipate that 2022 will include further insulation and modification of C2 servers to avoid detection. In response to 2020 takedown attempts, TrickBot's operators hardened their infrastructure to include more global VPS servers, in contrast to their previous reliance on compromised MikroTik routers. We anticipate the lessons learned from action against TrickBot and Emotet will lead to increased reliance on compromised devices, regular recycling of infrastructure, and the use of more resilient traffic encryption methods.

We also believe that similar methods will be employed by Cobalt Strike users, to protect Team Servers from the prying eyes of researchers. This will likely result in the dropping of traffic from known scanning engines and the use of redirects to mask the location of the Team Server (or other C2 nodes). This trend has likely already begun; BlackBerry's findings based on extracting configurations and internet scanning, identified roughly twice as many Team Servers as scanning alone. As internet scanning has become table stakes for intelligence providers, adversaries and red team operators have taken notice. We believe this is due to improved insulation of the Team Servers.

Finally, we expect the C2 environment to continue to diversify. As new malware families and C2 frameworks are released, we anticipate a portion of them will be aware of threat intelligence measures to scan and detect their servers. This will likely lead to partially decreased efficacy for threat intelligence scanning efforts in the short term but will result in new creative methodologies in the medium to long term. The detection cat-and-mouse game will continue, with both sides continuing to innovate to thwart the other.

Mitigations and Recommendations

To help safeguard systems, we advise the following mitigations:

- Use the Recorded Future Platform to help identify actively exploited vulnerabilities and CVEs that have been positively associated with ransomware variants, which can help with patch management and prioritization.
- Keep systems and software up to date and have a reliable and tested backup method.
- Exposed RDP servers are abused by threat actors to gain initial access into a target's network. If remote access solutions are crucial to daily operations, all remote access services (such as Citrix or RDP) should be implemented with multi-factor authentication.
- Password-protect sensitive files using strong, complex passwords.
- Detection in Depth
- Proactive detection creates an advantage for defenders, giving them preparatory time to ensure additional file- and network-based detections are in place.
- Employ detection-in-depth for common open-source tooling via correlation searches and Sigma queries in SIEMs for suspicious behaviors, YARA for suspicious file contents, and Snort for suspicious or malicious network traffic.
- The detections for each family show the increased use of open-source tools beyond just the families that get major publicity. Prioritize these families for network and host-based detection in enterprise environments.
- External network detections are only part of the detection equation; detection-in-depth methodologies, such as calculating the standard deviation of beaconing intervals or using YARA for memory inspection, can aid in the identification of malicious activity.

About Insikt Group®

Recorded Future's Insikt Group, the company's threat research division, comprises analysts and security researchers with deep government, law enforcement, military, and intelligence agency experience. Their mission is to produce intelligence that reduces risk for clients, enables tangible outcomes, and prevents business disruption.

About Recorded Future®

Recorded Future is the world's largest provider of intelligence for enterprise security. By combining persistent and pervasive automated data collection and analytics with human analysis, Recorded Future delivers intelligence that is timely, accurate, and actionable. In a world of ever-increasing chaos and uncertainty, Recorded Future empowers organizations with the visibility they need to identify and detect threats faster; take proactive action to disrupt adversaries; and protect their people, systems, and assets, so business can be conducted with confidence. Recorded Future is trusted by more than 1,000 businesses and government organizations around the world.

Learn more at recordedfuture.com and follow us on Twitter at @RecordedFuture.