Poster: Mining Threat Intelligence from Billion-scale SSH Brute-Force Attacks

Yuming Wu*§, Phuong Cao*§, Alexander Withers†, Zbigniew T. Kalbarczyk* and Ravishankar K. Iyer*
* University of Illinois at Urbana-Champaign,
† National Center for Supercomputing Applications

***Abstract*—This paper presents a longitudinal study of 11 Billion SSH brute-force attacks targeting an operational system at the National Center for Supercomputing Applications. We report the nature of these attacks in terms of i) targeted strategies (i.e., using stolen SSH keys), ii) large-scale evasion techniques (i.e., using randomized SSH client versions) to bypass signature detectors, and iii) behaviors of human-supervised botnet.**

The significance of our analyses for security operators include i) discerning cross-country attacks versus persistent attacks, ii) notifying cloud providers and IoT vendors regarding stolen SSH keys for them to verify the effectiveness of software patches, iii) deterring the above evasion techniques by using anomaly detectors/rate limiters, and iv) differentiating between fully automated attacks versus more sophisticated attacks driven by human.

I. Introduction

The Secure Shell (SSH) is the universal authentication protocol for managing remote servers. Attacks targeting exposed SSH servers see an exponential growth recently due to the availability of leaked passwords [9] and stolen keys [12]. A successful SSH login typically grants the super-user (root) permission, thus enables persistent access for compromising internal network, exfiltrating sensitive data [11] and causing monetary losses. For example, when being offered 50 Bitcoins by a hacker, a former server administrator at ShapeShift [7], a cryptocurrency company, gave away an SSH private key to the company's Bitcoin core server for accessing internal Bitcoin's wallets. This incident eventually led to $230,000 losses [1].

This paper presents a longitudinal study of 11 Billion SSH brute-force attacks targeting an operational system [6], [10] at the National Center for Supercomputing Applications[1] (NCSA). We report the nature of these attacks in terms of i) targeted strategies, ii) large-scale evasion techniques to bypass signature detectors, and iii) behaviors of human-supervised botnet.

A. Data Overview

Our dataset contains 11 billion attack attempts, including 3.4 billion connections and 7.9 billion SSH password- and key-based brute-force attack records. Each is an attempt to compromise the SSH server and thereby access the internal network and steal sensitive data. The data is collected in an operational honeypot in 1,000 days starting in February 2017, deployed on a `/16` IP address space simulating $\sim$65K machines [6], [10]. In total, the honeypot recorded 4.5 million unique, globally distributed, IP addresses of attackers.

B. Analysis Workflow

The main steps in our analyses are to: i) discern the nature of attacks in terms of persistence, ii) identify coordination and evasion techniques, and iii) distinguish human-supervised and fully automated botnet attacks.

II. Exploitation, Coordination, and Evasion

This section presents the exploitation, collaborating, and evasion strategies of advanced adversaries.

A. Exploitation of Leaked SSH Keys

In total, 185 unique SSH public key fingerprints (in the SHA-256 hash) found their way into our honeypot. By matching each of the keys with a public database and online files of bad keys [3], [5], we discovered and recovered the identities of seven keys that were publicly leaked due to vulnerabilities. Further investigations implied that cybercriminals were trying to gain root permission to vulnerable production appliances and devices in the wild using these leaked keys, even years after the key-pertinent vulnerability disclosure.

Attackers were targeting production devices using leaked keys. The seven leaked keys belonged to seven different enterprises. All these keys granted attackers with root permission in the targeted systems eventually.

The attackers used the privilege level related to each corresponding leaked key when targeting our honeypot. Instead of randomly using leaked keys to brute-force, the attackers have adequate details about pertinent vulnerabilities when plotting the targeted attacks.

Attackers were rapidly exploiting the leaked keys. Attacks that originated from Google LLC (Google), Charter Communications, and Portlane participated in exploiting the seven leaked keys. In particular, attackers from Google tried all seven identified, leaked keys, together with three other unidentified keys, over two days (July 28-29, 2018).

B. Key-based Collaboration

We inspected the diversity of attack sources using SSH keys in general, from which we uncovered global coordination.

An SSH key was exploited by 20 countries. We sorted key-based attempts characterized by the number of the originating IP address. Each of the ten keys originated from more than 15 distinct IP addresses, with the highest number being 71. However, most attackers originated from just a single country or AS. The only exception was used by 64 IPs scattered over 20 countries from four continents (including Asia, Europe, North America, and Oceania).

A persistent, single-country botnet versus a rapid, globally colluding botnet. Further inspection revealed that this globally coordinated botnet exploited a single SSH key 90 times within only four days (Dec. 11 to Dec. 14, 2017). On the other hand, one key originated from 71 IPs, yet all from a single country and AS. In contrast with the global botnet, this botnet persistently used one key for 2,700 times spanning five months (Feb. 2017 to July 2017). Compared with this single-origin bot, we can conclude that the globally coordinated bot wrapped up its fruitless attacks and shifted targets 50$\times$ faster.

C. Client Version-based Coordination and Evasion

Starting from August 2018, the honeypot witnessed an unprecedented emergence of unseen client versions: Over 1.7 million new client versions sprang up in August, 8,000$\times$ more than the total number of unique client versions in previous 18 months. Further inspection revealed only several hundred IPs spoofed client versions by randomizing over one million

§Joint first authors

[1]NCSA hosts the Blue Waters, a sustained petascale supercomputer of 22,640 cores.

OpenSSH version banners. This is unusual because about 90% of all IPs advertised only one client version. We speculate these randomizations were the attackers' mimic technique responding to our honeypot's deceitful defense mechanism.

Attackers randomized SSH client version banners at high frequencies. The top-spoofing IP advertised 400,000 unique client versions during its 200-hour attack campaign, implying varying an average of 2,000 client versions per hour. This attacker advertised `SSH-2.0-OpenSSH_7` within first several days of attack, then switched to massive spoofing by appending `SSH-2.0-OpenSSH_` with 5-character random strings.

A globally-coordinated botnets were involved in forging a million permutations of client versions. Only several hundred IPs were involved in scheming the large-scale randomization, and over 85% of them were new-incoming IPs in August. Further investigation showed that less than 300 IPs, yet globally coordinated from over 30 countries across six continents (all excluding Antarctica), actually accounted for the million-scale random permutations of client versions to masquerade their at

Defense-targeting evasion. The honeypot deceives attackers by randomizing key fingerprints for each of the 65,536 servers on the entire `/16` IP address space. We therefore suspect that the attackers were mimicking our honeypot's defense mechanism. Besides, attackers were massively hiding essential attack signatures, which would generally invalidate signature-based detection. Therefore, it calls for deploying new defense strategies such as rate-limiting or anomaly-based detection.

III. Human-supervised attack techniques

After discovering routine patterns of human attackers on a weekly basis, we further provide case studies to compare and contrast the distinctive behavior patterns and strategies between fully automated botnets and human-supervised botnets.

A. Data-driven methodology

Current work [11], [13] implemented additional features to capture human-generated activities, e.g., keyboard/mouse typing/clicking, window resizing. However, these methods introduced overhead to networking system design. Instead of modifying or adding features to the current design, The billion-scale attack attempts motivate us to come up with a data-driven methodology for mining human activity patterns.

Tail analysis of attack distributions. We focused only on IPs originating from one time zone. Then we chose a month with the most attack attempts. After grouping by IP, we computed average weekday and weekend attempts for each IP during the selected month. To quantitatively capture routine human evidence, we calculated a ratio of a weekday to weekend average attempts for each IP. Since we aimed to find relatively long-term (4–6 weeks) evidence, we filtered out IPs with the number of active weekdays lower than 15. Specifically, we then focused on IPs with the ratio Z-score [8] greater than three standard deviations (3σ) from the mean (μ), the tail on the rightmost part of the distribution.

Activity patterns of the human-supervised botnet. It turned out that all IPs in the tail, with similar activity patterns, originated from the same `/8` subnet, indicating organized routine management of botnet by the human attacker(s) [14]. These IPs also used the same client versions, passwords, and usernames. The daily intensity of these bots indeed aligned with human social routine on a weekly basis: periodic variations with decreasing activities on weekends (especially Sundays).

B. Case studies of two botnet types: human-supervised and fully automated

We selected another IP with a weekday to weekend average attempt ratio equaling to one and offer detailed case studies to distinguish the attack strategies adopted by both botnet types.

Human-supervised botnet is more resourceful in terms of attack devices. All bots iterated over four client versions with equal distribution for each. There were cases when these four client versions were used at the same time by one bot. On the other hand, the fully automated bot advertised one and only one commonly-used client version. Therefore, human-supervised botnet employed a more diverse handful of devices to launch attacks.

Human-supervised botnet is more ambitious and strategic in terms of credential brute-forcing. We used Dropbox zxcvbn [15] to measure password strength. For a fully automated bot, only one password (`7ujMko0admin`) scores 3, which is the highest among all 42 unique passwords it attempted, with the majority scoring 0. On the other hand, around 3,000 passwords used by the human-supervised botnet score 4.

Notably, one password used by human-supervised botnet begins with `Branch:masterFindfileCopypath`, and ended with a path in a Github repository [2]. This Git repo contains a wide range of passwords collected from backdoors, web shells, mail servers, WebLogic, Linux OS, dictionaries, etc. In addition to passwords, we also found collections of database and backdoor file paths, plus a script for brute-forcing enterprise mail servers, including Exchange [4]. On the other hand, fully automated bot rotated all 42 passwords every day over the entire attack campaign. Most passwords are commonly-used default passwords in Linux OS, IoT devices, routers, and firewalls.

IV. Conclusion

We investigated a broad scope of attack strategies in billion-scale SSH brute-force attacks. We discover great potential in attackers to launch large-scale, persistent, and evasion attacks that are accompanied by strategic human supervision. Also, we contribute methodology to cluster bot collaboration campaign in the wild, offer a scientific data-driven approach to differentiate between human-supervised versus fully automated botnet.

References

[1] "Shapeshift lost $230k in string of thefts, report finds - coindesk," 2016, https://www.coindesk.com/digital-currency-exchange-shapeshift-says-lost-230k-3-separate-hacks.

[2] "Blasting_dictionary," 2017, https://github.com/rootphantomer/Blasting_dictionary.

[3] "Ssh bad keys," 2017, https://github.com/rapid7/ssh-badkeys.

[4] "Enterprise email service for business - ms exchange email," 2019, https://products.office.com/en-us/exchange/email.

[5] "Packet storm," 2019, https://packetstormsecurity.com/.

[6] "Sdaia," 2019, https://wiki.ncsa.illinois.edu/display/cybersec/SDAIA. [Online]. Available: https://wiki.ncsa.illinois.edu/display/cybersec/SDAIA

[7] "Shapeshift," https://shapeshift.io/#/coins, 2019.

[8] "Z-score definition," 2019, https://www.investopedia.com/terms/z/zscore.asp.

[9] P. Cao, H. Li, K. Nahrstedt, Z. Kalbarczyk, R. Iyer, and A. J. Slagell, "Personalized password guessing: a new security threat," in *Proceedings of the 2014 Symposium and Bootcamp on the Science of Security*, 2014, pp. 1–2.

[10] P. M. Cao, Y. Wu, S. S. Banerjee, J. Azoff, A. Withers, Z. T. Kalbarczyk, and R. K. Iyer, "CAUDIT: Continuous auditing of SSH servers to mitigate brute-force attacks," in *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. Boston, MA: USENIX Association, Feb. 2019, pp. 667–682. [Online]. Available: https://www.usenix.org/conference/nsdi19/presentation/cao

[11] F. Dang, Z. Li, Y. Liu, E. Zhai, Q. A. Chen, T. Xu, Y. Chen, and J. Yang, "Understanding fileless attacks on linux-based iot devices with honeycloud," in *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 2019, pp. 482–493.

[12] V. Ghiëtte, H. Griffioen, and C. Doerr, "Fingerprinting tooling used for {SSH} compromisation attempts," in *22nd International Symposium on Research in Attacks, Intrusions and Defenses ({RAID} 2019)*, 2019, pp. 61–71.

[13] R. Gummadi, H. Balakrishnan, P. Maniatis, and S. Ratnasamy, "Not-a-bot: Improving service availability in the face of botnet attacks." in *NSDI*, vol. 9, 2009, pp. 307–320.

[14] S. Udhani, A. Withers, and M. Bashir, "Human vs bots: Detecting human attacks in a honeypot environment," in *2019 7th International Symposium on Digital Forensics and Security (ISDFS)*. IEEE, 2019, pp. 1–6.

[15] D. Wheeler, "zxcvbn: Realistic password strength estimation," *Dropbox Tech Blog, Apr*, 2012.

Mining Threat Intelligence from Billion-scale SSH Brute-Force Attacks

Yuming Wu[1 §], Phuong M. Cao[1 §], Alex Withers[2], Zbigniew T. Kalbarczyk[1], Ravishankar K. Iyer[1]
[1]University of Illinois at Urbana-Champaign, [2]National Center for Supercomputing Applications

INTRODUCTION/BACKGROUND

a longitudinal study of 11 Billion SSH brute-force attacks targeting an operational system at the National Center for Supercomputing Applications. We report the nature of these attacks in terms of

- persistence (i.e., consecutively attacking over an entire year)
- targeted strategies (i.e., using stolen SSH keys
- large-scale evasion techniques (i.e., using randomized SSH client versions) to bypass signature detectors
- behaviors of human-supervised botnet

The significance of our analyses for security operators include

- discerning cross-country attacks versus persistent attacks
- notifying cloud providers and IoT vendors regarding stolen SSH keys for them to verify the effectiveness of software patches
- deterring the above evasion techniques by using anomaly detectors/rate limiters
- differentiating between fully automated attacks versus more sophisticated attacks driven by human

LONGITUDINAL PERSPECTIVE OF ATTACK BEHAVIORS

Trend Anomalies

- Abrupt upsurge of unseen client versions from new attackers

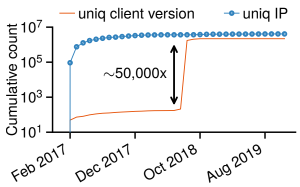

- Increasing scale of attack attempts from fewer attackers

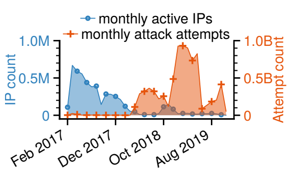

Persistent Attack Traces

- Persistent attackers constituted over 70% of all attacks

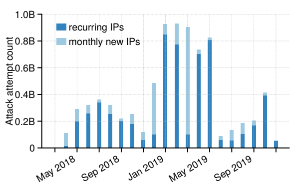

- Persistent attackers continuously attacked for over one year

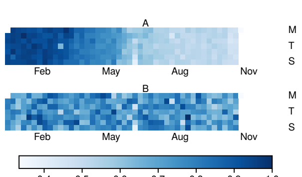

ANALYSIS WORKFLOW

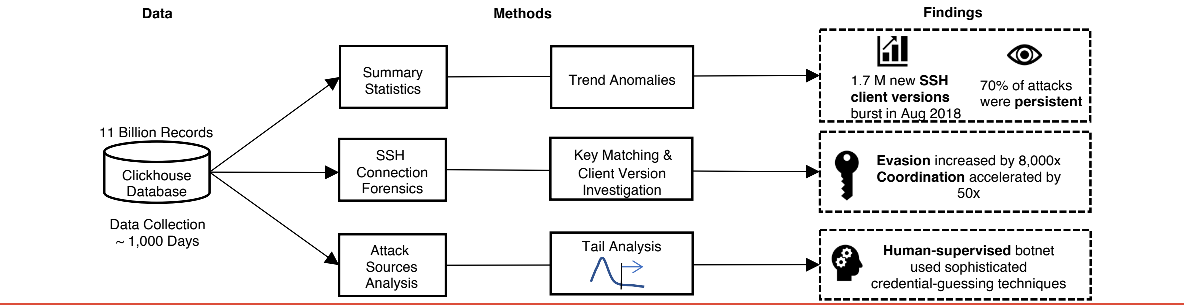

EXPLOITATION, COORDINATION, AND EVASION

Exploitation of Leaked SSH Keys

We discovered and recovered the identities of seven keys that were publicly leaked due to vulnerabilities.

- Attackers were targeting production devices using leaked keys

SSH Key (SHA256)	Key Owner	Appliance Type	Public Disclosure Year	1st Attack Attempt Year	Username
1M4Rz...qu0ZA	Vagrant [1]	Base box for development environments	2010	2018	root
9prMb...Ghro4	F5 [2]	BigIP appliances	2012		
MEc4H...UfTww	Loadbalancer [3]	Virtual load balancer	2014		
VtjqZ...PiQPc	Quantum [5]	Virtual deduplication backup appliance			
/JLp6...POCc0	Array Networks [4]	Virtual application delivery controllers Secure access gateways			sync
Z+q4X...8kIxM	Ceragon [6]	IP traffic router	2015		mateidu
f+1oG...zEDhc	VMware [12]	Data Protection appliances	2016		admin

- Attackers were rapidly exploiting the leaked keys

Attackers from Google tried all seven identified, leaked keys on the same day (Dec 14, 2018).

Autonomous System	Client Version [SSH-2.0-]	SSH Key (SHA256) & Key Owner						
		1M4Rz... Vagrant	9prMb... F5	MEc4H... Loadbalancer	VtjqZ... Quantum	/JLp6... Array Networks	Z+q4X... Ceragon	f+1oG... VMware
Google LLC	libssh_0.7.0	✓	✓	✓	✓	✓	✓	✓
Charter Communications	Ruby/Net::SSH...		✓	✓	✓	✓	✓	✓
Portlane	libssh-0.6.1			✓	✓			

Ruby/Net::SSH...refers to Ruby/Net::SSH_5.0.2 x86_64-linux-gnu.

Key-based Collaborations

- An SSH key was exploited by 20 countries
- A persistent, single-country botnet versus a rapid, globally colluding botnet

SSH Key (SHA256)	# Countr(y/ies)	# AS(es)	# IPs	Client Version [SSH-2.0-]
qlIN/...	**20**	**38**	**64**	**Go**
B6kr4...		2	25	
mumiE...	1	1	49	libssh-0.5.2
jSCqa...	1	1	42	libssh-0.5.2
V6OOC...	1	1	28	libssh-0.5.2
zPA6Y...	1	1	23	libssh-0.5.2
NH5Y7...	1	1	19	libssh-0.5.2
OyHmn...	1	1	17	libssh-0.5.2
8blLD...	1	1	16	libssh-0.5.2
+UJNI...		1	71	kthrssh__x00

Client Version-based Coordination and Evasion

- Attackers randomized SSH client version banners at high frequencies
- A globally-coordinated botnets were involved in forging a million permutations of client versions

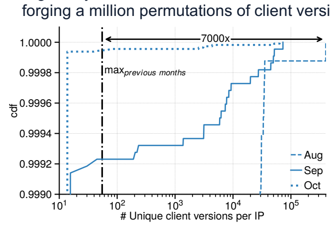

HUMAN-SUPERVISED ATTACK TECHNIQUES

Why understanding human-supervised attacks is important and our data-driven methodology

- Tail analysis of attack distributions
- Activity patterns of the human-supervised botnet

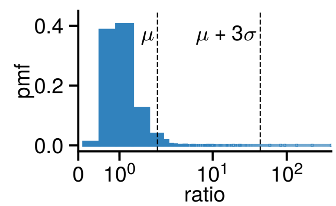

Case studies of two botnet types: human-supervised and fully automated

- Human-supervised botnet is more resourceful in terms of attack devices
- Human-supervised botnet is more ambitious and strategic in terms of credential brute-forcing

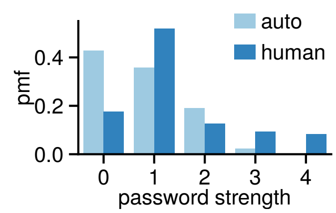

FINDINGS AND IMPLICATIONS

- Persistent attacks versus cross-country attacks
- Leaked SSH keys exploitation
- Large-scale evasion techniques
- Human-supervised botnets

HUMAN-SUPERVISED ATTACK TECHNIQUES (Cont.)

Type	Illustration of Daily Attempts [May 27 – July 21, 2019 (8 weeks)]	List of Unique Client Versions [SSH-2.0-]	List of Unique Username(s)	# Unique Passwords
Human-supervised	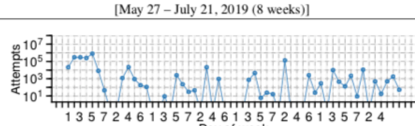 **Figure 8:** Human-supervised attack attempts declined on weekends [1: Mon. ~ 7: Sun.].	PuTTY OpenSSH_5.3 OpenSSH_6.2p2... nsssh2_4.0...	root	35,952
Fully automated	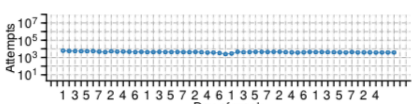 **Figure 9:** fully automated attack patterns were almost unvarying over different days of the week [1: Mon. ~ 7: Sun.].	sshlib-0.1	root, user, admin, ubnt, usuario, pi, supervisor, support, service, mother	42

OpenSSH_6.2p2...refers to OpenSSH_6.2p2 Ubuntu-6;
nsssh2_4.0...refers to nsssh2_4.0 NetSarang Computer, Inc.

ACKNOWLEDGEMENTS

SDAIA:
https://wiki.ncsa.illinois.edu/display/cybersec/SDAIA

NSF Grant: CICI: Secure Data Architecture: Shared Intelligence Platform for Protecting our National Cyberinfrastructure. Award Number: 1547249

DEPEND Symphony Cluster

§ Joint first authors

ILLINOIS