

CYBERATTACKS ON CLOUD HONEYPOTS

BY: MATT BODDY, SOPHOS

Contrary to popular belief, every device is worth hacking when the process is automated. It doesn't matter who or where you are, if you own a company big or small, or have technology in the home – every device can be monetized by an enterprising criminal. Brute force login attempts are likely occurring on any online device. Yet the speed and scale of the problem can boggle the mind. Criminals are relentless and often competitive with one another to find, take over, and monetize your smart devices.

The research you'll find here, using honeypot devices across the internet, is a first step in attempting to quantify the issue. In cybersecurity terms, a honeypot is an open, vulnerable device, configured to deliberately lure a cybercriminal to attack. When the criminal starts to interact with the device, they are in fact triggering alarms to alert a business or individual to their presence and track their activity.

There are many types of honeypots, but in this paper we focus on two main distinctions: high and low interaction.

A low-interaction honeypot is a honeypot that, once found by the hacker, will not be of much use to them. In our case, the attacker is presented with a login prompt they have no way of getting past. This logs and stores any attempts to log in, providing information on the attacker's IP address of origin (which can be attributed to a location), and the username and password used in the login attempt.

A high-interaction honeypot permits the attacker to go further in order to gather additional information about their intentions. In the context of this paper where high interaction honeypots are referenced, we allowed the attacker to log in to the honeypot with a designated set of usernames and passwords, and stored any command the attacker attempted to use.

The honeypots in this test simulate the Secure Shell (SSH) service and, therefore, measure SSH login attempts. SSH is a remote access service used not only by servers, but is also enabled in domestic environments in devices as diverse

as CCTV cameras or NAS devices. On these systems, legitimate users may connect via SSH to remotely configure the device or to access files. For an attacker, once they get past the login prompt onto an IoT device, they not only gain the same access as the owner, but often gain even more control than was ever intended.

We initially set up honeypots in ten of the most popular AWS data centers in the world and made sure that the honeypots are not affiliated with Sophos or any other company other than, perhaps, the hosting provider. To a hacker, they appear as just a number, a bit of extra processing power that could be theirs, a camera they could control or a directory of files they could access and share.

The research clearly demonstrates that devices that have not received due attention to configuration (including changing any default passwords installed at the factory on many devices) may permit a cybercriminal to access those devices. However, we can learn how attackers work from this research, and what we can do to prevent many of them from succeeding.

THE FINDINGS

Finding 1: The short time it takes to get pwned

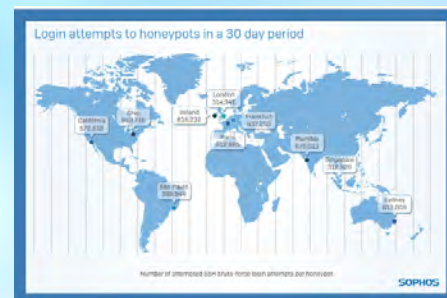
When the honeypots first went online, it took attackers no time at all to discover the SSH service and for login attempts to start. In one instance, our device was attacked in less than one minute from deployment. However, in others it took nearly two hours before login attempts began. But once the login attempts start, the attacks are relentless and continuous.

Finding 2: It is a feeding frenzy

Once the honeypots were well established, each device saw an average of 13 login attempts per minute, or about 757 per hour.

Finding 3: The Chinese connection

95.4% of the traffic we tracked appeared to originate in China. This doesn't necessarily mean that the attackers conducting these brute-force attempts are also located in China, because attacks may be routed through other machines under the attackers' control.



Finding 4: The global distribution of login attempts

The London honeypot alone suffered just over 314,000 login attempts over the course of the 30 days in which we ran these honeypots, with the honeypot hosted in Ireland suffering more than 600,000 login attempts. Other notable figures include over 950,000 attempts in Ohio.

Does this mean that hosting services in London is safer than hosting services in Ohio? In short, no. Honeypots based in every region received hundreds of thousands of login attempts over this 30 day period. These attempts varied in complexity from default usernames and passwords down to complex passwords with what security practitioners would consider sufficiently complex combinations of numbers, letters, and special characters. No one country is safer than any other. Wherever you are in the world, following good security practices is paramount.

DEFAULTS ARE THE FAULT

Looking at what drives this number of brute force login attempts, we found the dominant problem was ongoing exposure as a result of not changing default usernames and passwords.

For example, 'root' exists as a default username for most *NIX devices. Consequently, it is unsurprising that it is consistently at the top of the list of most seen username login attempts. However, the sheer scale is remarkable: 'root' accounts for 5,211,644 of the 5,447,956 logins (just under 96%). Because the 'root' account provides administrative access to devices, it's likely that, after their botnet reaches a significant size, the cybercriminal will use this privileged access to perform large

Username and potential associated devices

USERNAME	DEFAULT DEVICE	LOGIN ATTEMPT COUNT
root	Most Linux devices and many IoT devices (can be one and the same) including Raspbian, Synology NAS devices	5,213,494
admin	IoT devices including ACU, Access, AVTels, Basics, Beosystem, FLIR, GANZ, Hovalpro, Ictvision, Hikvision, Hunt Electronic, iCATCH, JVC, LG, Mobotix, Panasonic, Proxied and Streaming CCTV devices and Teleglobe, Verbatim and Lacie NAS devices	478,816
user	IoT devices	41,945
admin	Ubiquiti Networks' default username	3,962
ubuntu	ZTEK Ubuntu instance	3,300
hagene	Nagios network monitoring	3,200
pi	Default username on Raspbian	2,217
raspberrypi	Raspbian's default username	1,748

Source: Logon attempts from honeypots in 2018.

scale DDOS attacks to organizations and institutes as seen before in botnets like Mirai.

There are other correlations we can make between login attempts and specific technologies. For example, the username 'pi' was represented in the top 20 attempted usernames because it is the default username for Raspberry Pi-based computers running the Raspbian operating system. The fact that the username exists here shows that, through misconfiguration or negligence, these devices appear on the internet as exposed and vulnerable.

If a device is online, anyone can attempt logging in, so the only line of defense is the password. This is where we see opportunist cybercriminals are aiming for commonly used and poorly chosen passwords for the root account. In much smaller numbers we also notice login attempts targeting specific brands or models of device. These login attempts are focusing on devices still configured with well-known default accounts and passwords. For example, for Raspberry Pis running Raspbian, a distribution of Linux designed for the Raspberry Pi, the default password "raspberrypi" appears 1,808 times, taking it into the top 20 attempted passwords.

This dastardly duo of default login details and obvious links to hardware names gives hackers an easy ride. We have collated a list of the most recognizable default usernames and passwords that were used by attackers and, to illustrate the issue, we have suggested IoT device brands which use those defaults.

Going beyond the default password, there is a further issue with commonly chosen passwords. When analyzing a handful of the most seen password attempts, clear, simple keyboard patterns are being exploited by cybercriminals. Two canonical examples of this would be the passwords "1qaz-2wsx" and "1q2w3e4r." Looking at a keyboard, it is easy to see how all of these keys are located right next to each other.

Many devices ship with default passwords. Whether they are easy to guess or not, this is always a mistake. Any device with a default password quickly becomes widely known in the underground community and is no better than a "well known secret." Even when users are

forced to choose their own password during initial setup, we often observe little care is taken in choosing a secure password leading to easy attempts by criminals to brute force guess working combinations.

WHAT HAPPENS ONCE THEY'RE IN?

The purpose of this research was to establish the frequency, consistency, and complexity of the average attack on the average person. As a result, the majority of the research focused on low-interaction honeypots to measure the number of login attempts. However, the high-interaction honeypot was included to better understand what the average device may well be directed to do once compromised.

The research suggests that if you're unfortunate enough to have a weak username and password and your device ends up online, you will be involved in attacks aimed at large organizations.

From the high-interaction honeypot, we pulled this typical course of action:

1. Login attempt of username:root password:admin succeeded
2. TCP connection request to Yandex over HTTPS
3. TCP connection request to large retail chain's open API over HTTPS
4. TCP forward request to large retail chain's open API over HTTPS

The above process repeats thousands of times, making it appear automated. However, we can still analyze the steps in the attack.

1. Check that the honeypot has a valid internet connection by connecting to a well-known address. This is via a secure connection request to Yandex. Yandex is a popular search engine in eastern Europe and Russia.
2. The attack then checks if connectivity to the target service is available – in this case, a connection request to a remote IP address belonging to a large retail chain's open API.
3. There then follows an attempt to exploit large retail chain's IP address using the SSH honeypot server as a proxy.

By being compromised, the honeypot has now become an amplification device for the cybercriminal to launch further attacks on other infrastructure.

CONCLUSIONS AND COUNSEL

In light of the above observations, we have some initial recommendations to keep devices secure and break these botnet chains. Most of the login attempts preyed on default usernames and/or passwords. Changing these is a critical initial step to improving the security profile of a business, and it must be applied rigorously to all new

Passwords and potential associated devices

PASSWORD	DEFAULT DEVICE	LOGIN ATTEMPT COUNT
123456	IoT devices including ACU iCATCH and Sea Max CCTV cameras	15,778
admin	IoT devices including Acas, AVTels, Basics, Beosystem, Goveon, Hunt Electronic, Kantix, iCATCH, LG, Proxied, Sanyo and Streaming CCTV cameras and Synology and Lacie NAS devices	12,806
1234	IoT devices including GANZ ProxiedPro CCTV camera	9,663
password	IoT devices including Digiscam modules, Small Logic NAS devices	8,029
12345	IoT devices including Ictvision and Panasonic CCTV camera	7,336
user	Ubiquiti Networks' default password	6,217
root	IoT devices including devices manufactured by 50+ IoT CCTV cameras	5,707
123	IoT devices including Teleglobe CCTV cameras	5,433
blank	The blank password is used on Axis and Videauk IoT devices as the default password	3,242
raspberrypi	Raspbian is the default password for Raspbian, the Raspbian's distribution of Linux	1,808



devices. The recommendation is simple – change all passwords from the default and avoid obvious patterns.

There is also a specific weakness around universal plug and play (UPnP). UPnP automatically sets up a port forwarding rule which allows connections between routers and devices. Although the research has not addressed this specifically, it could be the reason for so many login attempts targeting CCTV and other IoT devices. The simple advice is to turn off UPnP on routers.

Ideally, everyone should use complex and unique passwords for each service. To make this simple, they should use a password manager. A password manager can introduce unique complex passwords for every website you use, and you only have to remember a single password for the password manager itself. And on SSH servers, use key based authentication, not just a password. Key-based authentication provides an alternative to password based authentication; if you don't have the key, you're not allowed in. Where available, administrators can deploy tools such as fail2ban on Linux servers to limit the number of login attempts someone can make before their IP address is banned from connecting again.

If an attacker does manage to get onto your Linux device, then your last line of defense should be a malware scanner, such as Sophos Antivirus for Linux, which will catch known payloads that are dropped onto your device by the adversary.

The aggressive speed and scale of attacks on new devices should send a strong signal to anyone working with technology, as well as deploying technology around the home. The only default setting that is acceptable is that of caution and best practice. It is not a case of if you will be targeted, but rather when, and how prepared you are for the attacks to come. **SME**