Cloud-as-a-Cyber Weapon
Trial Accounts and Account Takeover Arming Cyber Criminals

An Excerpt

Stratecast Perspectives & Insight for Executives (SPIE) Vol. 19 no. 7

Michael Suby
Stratecast VP of Research

Stratecast | FROST & SULLIVAN

50 Years of Growth, Innovation and Leadership

April 2019

INTRODUCTION[1]

How Distributed Denial of Service (DDoS) attacks are perpetrated is continuously evolving. In 2016, Internet of Things (IoT) DDoS attacks gained notoriety with the Mirai IoT botnet.[2] As botnet DDoS attacks were peaking, cyber criminals were refining other attack methods. In 2018, large-scale amplification DDoS attacks capitalizing on internet-exposed memcached servers came to the forefront.[3] The 1.35 terabits per second attack on GitHub is a prominent example.[4]

In comparison to botnet-enabled DDoS attacks, memcached amplification attacks do not require building an army of malware-infected IoT devices to stage an attack. Instead the attackers' staging footprint is much smaller. Using IP address spoofing techniques, the attacker sends a series of small requests identified with the victim's IP address to multiple internet-connected memcached servers. DDoS traffic amplification occurs as the memcached servers send back much larger responses (amplified by 50x) to the victim's genuine IP address for processing. The intended outcome is the same as botnet attacks: inundate the victim's web services with more processing demands or volume of requests than serviceable, such that legitimate requests are either blocked or responses are delayed to the point of uselessness.

Depending on the target and the size and duration of the attack campaign, the business implications can be significant. For GitHub, it suffered 10 minutes of intermittent outages until Akamai was called in to scrub GitHub's traffic free of attack traffic. Eight minutes later, with the attack neutralized, the attacker retreated. Although service was fully restored, GitHub sustained a financial cost of paying Akamai to fight the attack, in addition to the business impact of the intermittent outage. In another high-volume DDoS attack involving more than 700,000 HTTP requests per second, GreatFire.org stated that the attack cost it $30,000 per day in AWS hosting charges to scale its service-supporting infrastructure until the attack subsided.[5]

Clearly, DDoS attacks are not victimless crimes. The owners of the attacked services are victims, but so are legitimate users of the disrupted services. For those users, their service experiences are diminished and, depending on the service's purpose and user significance, users, particularly B2B relationships, can suffer tangible loss. Unfortunately, further evolution in how DDoS attacks are perpetuated will snare additional victims.

Discussed in this insight is the fraudulent use of public cloud trial accounts (also known as synthetic accounts) and public cloud account takeover (ATO) by DDoS attackers to wage their attacks. Similar to DDoS attacks that capitalized on an IoT botnet or memcached servers, without paying for those assets, public cloud service assets (e.g., Microsoft Azure, AWS, Alibaba Cloud, and Google Cloud Platform) are being fraudulently capitalized by DDoS attackers. Intuitively, the owners of these assets, the public cloud providers, are victims. Cloud customers, as Stratecast describes in this insight, are also victims.

Additionally in this insight, Stratecast describes how Arkose Labs "attacks" the economic model of account fraud in public clouds.

[1] In preparing this report, Stratecast conducted interviews with:
- Arkose Labs – Kevin Gosschalk, Co-Founder & CEO

Please note that the insights and opinions expressed in this assessment are those of Stratecast, and have been developed through the Stratecast research and analysis process. These expressed insights and opinions do not necessarily reflect the views of the company executives interviewed.

[2] *Inside the infamous Mirai IoT Botnet: A Retrospective Analysis* (December 2017) describes how Mirai IoT botnet attacks operate.

[3] *Memcrashed - Major amplification attacks from UDP port 11211* (February 2018) describes how amplification DDoS attacks operate.

[4] *GITHUB SURVIVED THE BIGGEST DDOS ATTACK EVER RECORDED* (March 2018)

[5] *GitHub DDoS Attack Traces to China-Disruption Appears to Target Anti-Censorship Tools* (March 2018)

THE RISE OF DDOS ATTACKS EMANATING FROM PUBLIC CLOUDS

With each public cloud assigned a unique Autonomous System Number (ASN), tracking traffic by ASN is possible. Link11, a cloud-based anti-DDoS protection service provider, has assembled data for ASNs assigned to public cloud platforms. A snapshot of that data in Figure 1 shows the rise in DDoS attacks from public clouds. Other providers of DDoS mitigation services are also gathering evidence on this same trend but have not yet released their findings.

Figure 1: DDoS Attacks Shifting Towards the Cloud

In the 12 months ending June 2018, 25% of DDoS attacks in Europe used abused cloud servers; an increase from 18.5% in the previous 12 months.

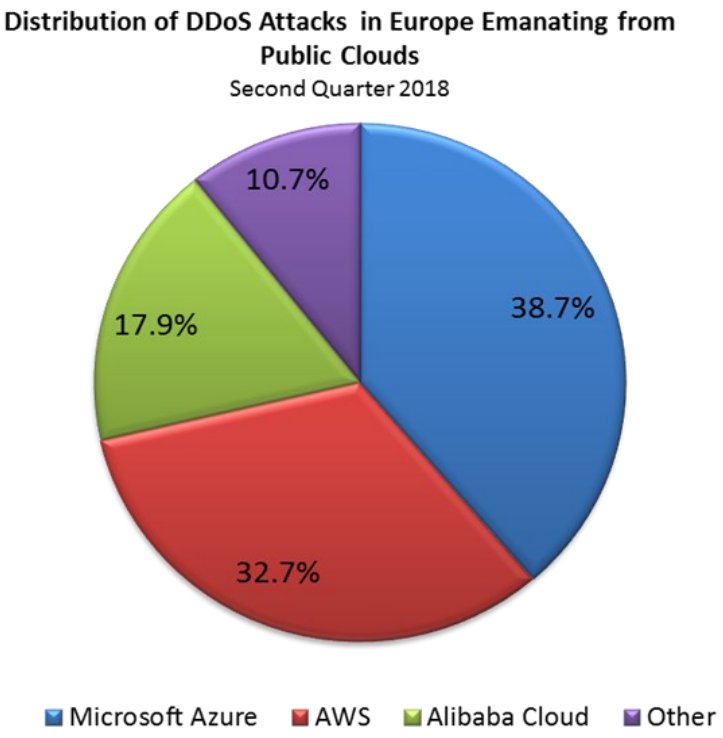

Source: Analysis from Link11's Security Operations Center

WHY PUBLIC CLOUDS ARE USED FOR DDOS ATTACKS

The reasons why DDoS attackers view public cloud accounts as an attractive environment to stage their attacks are quite logical. In this section, we present those reasons.

DDoS attackers' business model produces attractive returns with free resources

As DDoS attacks are dependent on sending packets/requests to targeted web services, compute and networking resources are essential ingredients. There is no better way to gain a favorable return on effort then when the cost of these resources is small or free, as they are with trial accounts or clandestine use of customer accounts (i.e., ATO).

Turning cloud accounts into launch pads for DDoS attacks is in cybercriminals' wheelhouse

All public cloud providers offer free trial accounts, as the freemium model introduces non-cloud users to the power of cloud services without a financial burden. With a favorable experience, trial users are apt to upgrade as commercial customers.

To further build user attraction, cloud providers offer a range of services in trial accounts, as evident in the trial account webpages of Microsoft Azure and AWS (click on each to view). What appeals to legitimate users also appeals to cyber criminals. An extensive set of cloud services provides a broad toolset for facilitating a variety of attack scenarios. Standing up memcached servers in multiple trial accounts, across multiple cloud providers, is one plausible scenario.

Also evident in these webpages, account registration is required. For experienced cyber criminals, this is not an insurmountable task. They are masters in creating phantom digital identities and harvesting stolen identities. Also in the industrialization of their tradecraft, bots are likely used to create fake accounts with push-button efficiency. Going one plausible step further, dark web offerings that create and sell trial accounts, or DDoS-attacks-for-hire services built on trial accounts, are other insidious means that arm DDoS attacks.

The cloud providers do have procedures to prevent fake trial accounts, as noted in this passage from Microsoft Azure's Free Account FAQs page: *Why do I need to provide a credit card and phone number? One of the ways we keep prices low is to verify that account holders are real people, not bots or anonymous trouble-makers. We use the phone number and credit card for identity verification.* The rise in DDoS attacks emanating from public clouds does question this procedure's effectiveness.

Similar fraud approaches in trial accounts can play out in ATO. Cyber criminals leverage stolen credentials, or guess credentials (not too difficult if a password is simple in design, frequently reused, seldom changed, or shared with others) to crack into legitimate cloud accounts. With time-tested industrialized techniques used to crack into other types of accounts, such as accounts of web-based email services, DDoS attackers are likely either experienced or can hire the services of a criminal third party to crack into legitimate cloud accounts when access is controlled by only a username and password combination.

Barriers to success are low

With industrialized and commercialized cybercriminal tradecraft likely employed to create trial cloud accounts, or to commandeer legitimate accounts, the "how to" barrier in using public cloud accounts for DDoS attacks is marginalized. But this is only one barrier to overcome on the pathway to criminal success. Other barriers are concealing fraudulent use and attack blocking. Unfortunately, these barriers are unlikely to be significant, as we explain below.

Concealing Fraudulent Use

Although logical that concealment is a necessary element for cybercriminals, concealment is also dependent on actual policing on how an account is used. There are reasons to conclude that policing is either non-existent or weakly performed.

On non-existence, the cloud providers offer trial and legitimate accounts without strings attached. Users can consume available cloud services as they wish, for whatever purpose they decide. The cloud providers are not dictating usage or intent, as evident in this passage from Microsoft Azure's Free Account FAQs page: *Can the*

Azure free account be used for production or only for development? The Azure free account provides access to all Azure products and does not block customers from building their ideas into production. Moreover, peering into what is occurring within a trial or legitimate account would violate customer privacy. And, if the outbound traffic emanating from the cloud is encrypted, the cloud provider's visibility is obviously restricted, as are the actions it can take.

Commercial account owners have a greater financial obligation to self-police. The more cloud services consumed, the more they pay. It is in their fiduciary self-interest to monitor. But monitoring is not a guarantee of abuse detection. If, for instance, account review is conducted monthly, or even weekly, the cybercriminal may have already used the account for its nefarious purposes and moved on. In another "for instance," accounts can have multiple legitimate users. To detect fraudulent usage, the account owner would need to conduct an investigation to confirm if a spike in usage is genuine abuse, before suspending service consumption. Time to conduct the investigation is time the cybercriminal can operate undeterred.

One last point, sophisticated cybercriminals, particularly those that offer DDoS attacks as a service, recognize the detection potential, and will take steps to alleviate the risk of detection. Limiting or spreading out usage over time or over multiple accounts are approaches they will likely use.

DDoS Attack Blocking

Web service owners under DDoS attacks emanating from public clouds are in a precarious situation. The ASN associated with the DDoS attack traffic could also be associated with legitimate cloud-hosted services that the owner's business relies on. Blocking traffic by ASN would interfere with legitimate business operations—both bad and good traffic are blocked. From the attacker's perspective, blocking by ASN is a less likely scenario than if the attack traffic was emanating from a non-cloud ASN, which can be blocked without negative business consequences. Perversely, DDoS attackers are rewarded with less blocking risk by firing their attack traffic from public clouds.

EXPANDING VICTIM LIST

Owners of web services and their users are unquestionably victims of DDoS attacks. Unfortunately, with the fraudulent use of public cloud accounts for staging DDoS attacks, the victim list expands to include public cloud providers and their customers.

Directly, public cloud providers are incurring an uncompensated expense in the use of their cloud resources. Although trial accounts are complimentary by design, the uncompensated expense is the absence of account upgrades. The cloud provider, in essence, does not have a compensation avenue for this free promotional service as cybercriminals will not be upgrading their fraudulent trial accounts.

Stratecast believes that the public cloud providers are likely very hesitant to turn down their trial account programs. To do so would interfere with a mainstay program designed to attract new users; introduce new services and features, and accelerate user awareness; and compete effectively in the public cloud market.

Alternatively, public cloud providers could add more safeguards to further prevent fraudulent trial account creation. Doing so, however, will require additional investments by cloud providers; and worse, potentially add friction to the user experience of trial accounts. The latter is a penalty waged on the many legitimate trial account users due to the fraudulent behaviors of a few. The Arkose Labs system of fraud prevention, as discussed in the next section, is designed to only add painful friction to fraudulent trial account creators.

In ATO, commercial account owners are paying the price of illegitimate consumption of cloud services posted in their accounts. Although some account owners will make the case to the cloud providers to reduce their account fees for fraudulent usage, the account owners are saddled with the burden of proof and rallying the effort to make the fee-forgiveness case. Neither of these is totally expense-free to the account owner, and there is no guarantee of full forgiveness. Alternatively, the cloud providers could credit account owners for their claimed illegitimate fees, and possibly grant an extra credit for the inconvenience. In doing so, the cloud provider is negatively impacting its profit margin, as no remuneration is received for actual services rendered.

More theoretical, but still plausible, the public cloud providers could encounter backlash from governmental agencies. For this scenario, consider social media platforms. Maybe naively, they were designed with the intent to only foster beneficial outcomes. Even so, unintended and perhaps unpredictable negative outcomes have transpired. Public scrutiny intensified, and the extent of long-term business impact on social media platforms is yet to be fully determined. At minimum, the additional armies of human reviewers and data scientists tasked with purging these platforms of reprehensible content have surely increased the operating costs of social media platforms. Possibly the same sequence of events will occur with public cloud providers. If their platforms are deemed responsible for facilitating criminal cyber activity, with an unequal effort by the platform owners to prevent it, governmental agencies could step in with the intent of standing up for victims and taking a stand against cybercriminals.

UPENDING THE ECONOMIC MODEL OF ONLINE FRAUDSTERS

Arkose Labs has developed a distinctive approach to preventing online fraud. Recognizing that the economics of online fraud is tied to automation (i.e., bots), Arkose Labs's technology is designed to render automated online fraud uneconomical at the front door of account creation and account access. Another design principle is zero friction for legitimate users. Optimally, and as previously stated, legitimate users should not be penalized for the actions of fraudsters.

Arkose Labs is not a typical cybersecurity technologist. Its heritage is in electronic gaming and, machine learning and visioning. Also, the company does not subscribe to cybersecurity's typical multi-layered approach. Instead, the company set its sails on complete prevention rather than a combination of prevention and then mitigation to compensate for prevention's inadequacies. With complete prevention as its objective, Arkose Labs' standard customer agreement includes a service level agreement (SLA) that 100% of automated abuse will be prevented from achieving scale.

Shown below in Figure 2 are the two pillars that define Arkose Labs' technology: Telemetry and Enforcement.

Figure 2: Arkose Labs' Synergistic Telemetry and Enforcement Pillars

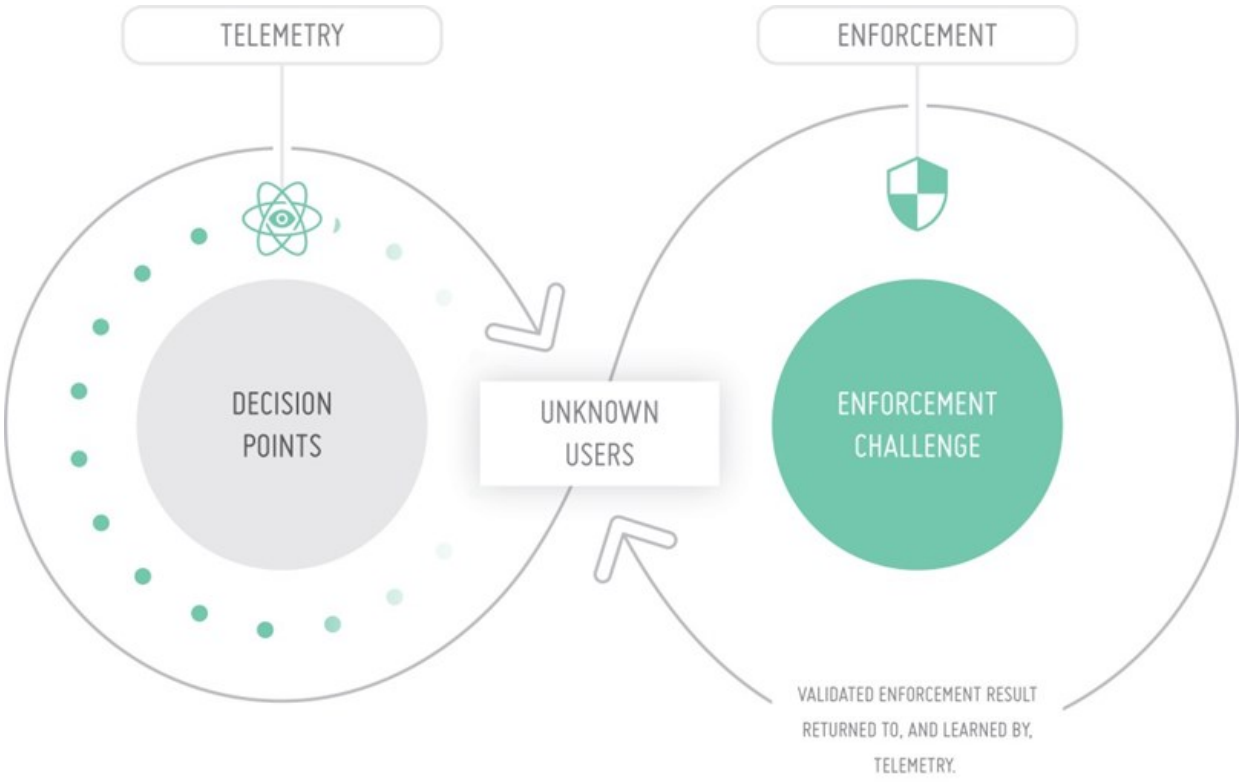

Source: Arkose Labs

Telemetry is the engine for determining whether a requestor is a human. Yet, despite advances in decision-making models and expansion in collectible data, telemetry as practiced by most online service companies has a grey area of uncertainty. In this grey area are the unknown users. For these users, they are neither certain humans nor certain machines. There is sufficient evidence to flag them as potential machines, but not enough to be 100% certain. Faced with this grey area of unknown users, online service companies are faced with several undesirable choices:

- **Allow unknown users to carry on** – With this choice, online fraud is solely addressed by the effectiveness of the telemetry model. The online service company, in essence, is making a conscious decision to accept some fraud in exchange for not inconveniencing legitimate users that are categorized as unknown. The fraud risk, however, is not static. It has the potential to increase as cybercriminals learn and adapt to the model's shortcomings.
- **Interrupt unknown users, and intervene with human decision-makers** – In this choice, humans are a second branch on a decision tree. Failing to determine with certainty if a requestor is human or machine, back-end humans make a second and hopefully more certain decision on user legitimacy. Not without consequences, the company incurs additional expense in a human-dependent decision process that is neither scalable nor infallible.

- **Interrupt unknown users with an online test** – This choice involves presenting unknown users with a challenge-response test, such as interpret and type the alphanumeric characters presented online or choose among a patchwork of online images that match a theme (e.g., which images are storefronts?). Effective to a degree, challenge-response tests are subject to being circumvented by cybercriminals as they advance their test-passing automation. These tests are also difficult to balance between being too complex for automated means to solve but sufficiently simple for humans to pass with inconsequential effort on first attempt. Described next, Arkose Labs's second pillar, Enforcement, creates balance between machine complexity and human simplicity. Also, with a real-time feedback loop, Enforcement also systematically improves Telemetry (becomes smarter), which contributes to a decline in the relative percent of unknown users (fewer subject to an "Are you a human?" test).

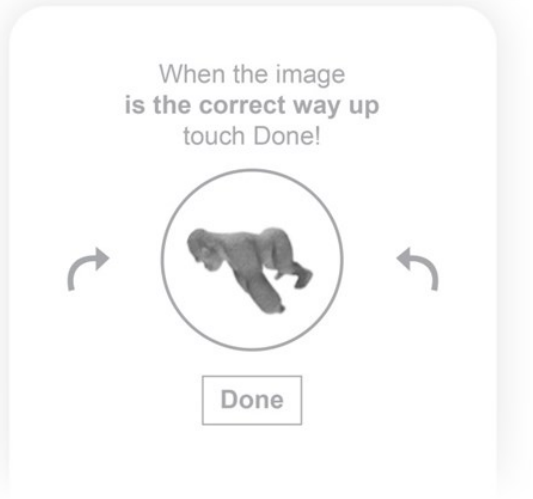

Patent-pending is Arkose Labs's 3D Enforcement, which presents unknown users with a question (i.e., challenge) to determine their authenticity based on the accuracy of the provided answer. In one such challenge, unknown users are presented with a single-use moveable 3D image (e.g., an image of a gorilla). The user clicks on directional arrows, and then "done," once the image is correctly oriented. While a simple task for humans, the same is not true for machines, as both the challenge (e.g., correctly orientate an image) and the 3D image (e.g., a gorilla) are variables that change with each instance. The corresponding "challenge" encountered by fraudsters in attempting machine automation in solving one challenge does not support solving future challenges—the future is not predictable from the past. With an extensive library of 3D images and challenge types, Arkose Labs has an infinite number of challenge and image permutations to stump fraudsters.

With Enforcement being highly accurate (i.e., only human solvable), Arkose Labs uses these results from across its base of customers and online fraud use cases (e.g., ATO, fake users, transaction blocking, and game hacking) to improve the Telemetry model. The result is that fewer users are subject to Enforcement. According to Arkose Labs, challenged users consistently represent a single digit percentage of legitimate users.

In closing, Arkose Labs offers an appealing approach to public cloud providers in combatting abuse of trial accounts and ATO. Customer experience demonstrates that Arkose Labs is effective in blocking automated abuse; does not impair the experience of legitimate users; and is easy to implement, as the Arkose Labs cloud-based technology engages users before they enter account creation or account access web portals.

THE LAST WORD

DDoS attacks emanating from fraudulent trial accounts or via ATO are likely not the only illegitimate use of public clouds. Cryptomining, spamming, phishing, command & control, malware delivery, credential harvesting, and phony eCommerce storefronts are other potential uses. Essentially any digital operation requiring compute and storage resources and big-bandwidth access to and from the internet is plausible. Fold in the instant scalability and stand up/tear down attributes epitomized by public cloud infrastructure services, and public clouds appeal to the "shell game" quality that benefits criminals. Movement from one account to another makes tracking down the perpetrators a challenging assignment.

Solving for fraudulent trial accounts is a problem for the public cloud providers to address. The vehicle they created to promote their wares is being used to steal their wares. Also, left undeterred, the public cloud providers open themselves up to scrutiny for not being good citizens on the digital planet.

Michael Suby

VP of Research
Stratecast | Frost & Sullivan
mike.suby@frost.com

Stratecast | FROST & SULLIVAN

SILICON VALLEY | 3211 Scott Blvd, Santa Clara, CA 95054
Tel +1 650.475.4500 | Fax +1 650.475.1571

SAN ANTONIO | 7550 West Interstate 10, Suite 400, San Antonio, Texas 78229-5616
Tel +1 210.348.1000 | Fax +1 210.348.1003

LONDON | Floor 3 - Building 5, Chiswick Business Park, 566 Chiswick High Road, London W4 5YF
TEL +44 (0)20 8996 8500 | FAX +44 (0)20 8994 1389

ABOUT STRATECAST

Stratecast collaborates with our clients to reach smart business decisions in the rapidly evolving and hyper-competitive Information and Communications Technology markets. Leveraging a mix of action-oriented subscription research and customized consulting engagements, Stratecast delivers knowledge and perspective that is only attainable through years of real-world experience in an industry where customers are collaborators; today's partners are tomorrow's competitors; and agility and innovation are essential elements for success. Contact your Stratecast Account Executive to engage our experience to assist you in attaining your growth objectives.

ABOUT FROST & SULLIVAN

Frost & Sullivan, the Growth Partnership Company, works in collaboration with clients to leverage visionary innovation that addresses the global challenges and related growth opportunities that will make or break today's market participants. For more than 50 years, we have been developing growth strategies for the Global 1000, emerging businesses, the public sector and the investment community. Is your organization prepared for the next profound wave of industry convergence, disruptive technologies, increasing competitive intensity, Mega Trends, breakthrough best practices, changing customer dynamics and emerging economies?