tripadvisor®

Credential-Probing Attacks

TripAdvisor Security Guild

2017-08-17

Bill Langenberg

tripadvisor®

Topics

- What a credential-probing attack looks like
- How to detect it
- How to blunt an attack in progress
- Industry best-practices
- Key take-aways

Overview

TripAdvisor detected a credential-probing attack the evening of April 24, 2017, based on a large event spike within our registration member audit logging. We determined that the attacker was playing credentials from list(s) compiled from other compromised sites around the ‘net. This first attack was dubbed the “Trusov” attack given the owner of all involved subnets was “Trusov Ilya Igorevych.” Using the patterns established in the Trusov attack, we uncovered a few other attacks, occurring over 15 intervals.

We believe the intent of this probing activity was to validate the credentials to increase value for resale. All compromised accounts were re-secured. There was no indication of fraudulent activity, or a data breach within TripAdvisor.

tripadvisor®

Anatomy of a Probing Attack | Precursor

- The perpetrator conducts manual testing to understand our flows, the required format of requests, and expected responses
  - This activity is lost in the noise of normal site behavior
- There is often a sentinel account that the perp uses to verify that the attack configuration is still valid before opening the floodgates
- When the probing is detected, using the pattern or identifying the sentinel account to look for the original testing is our best chance to catch the criminal

tripadvisor®

Anatomy of a Probing Attack | Detection

- The attacks may be highly distributed, making them difficult to detect and block via DoS filters
  - Trusov featured > 100 subnets from around the world, including the U.S., with ~41 K possible IP addresses
- Once the probing attack is underway via a botnet, it can still be hard to detect an anomaly in the context of general traffic
  - Having specific signals for suspicious behavior is critical
  - The event rate is substantially higher if you know what to look for...

Anatomy of a Probing Attack | Detection

- Long-lived login events
  - TA_LOGIN
  - Looks fine

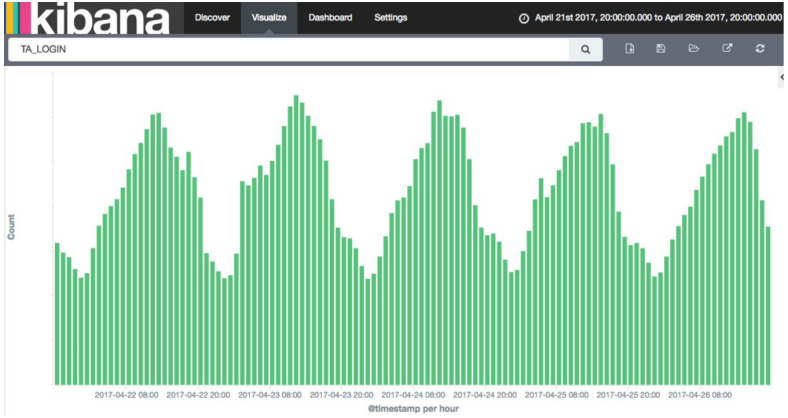

- Password update required
  - TA_UPDATE_PASSWORD_REQUIRED
  - Uh, oh!

- E-mail as the source
  - TA_CREDENTIALS
  - Hrm...

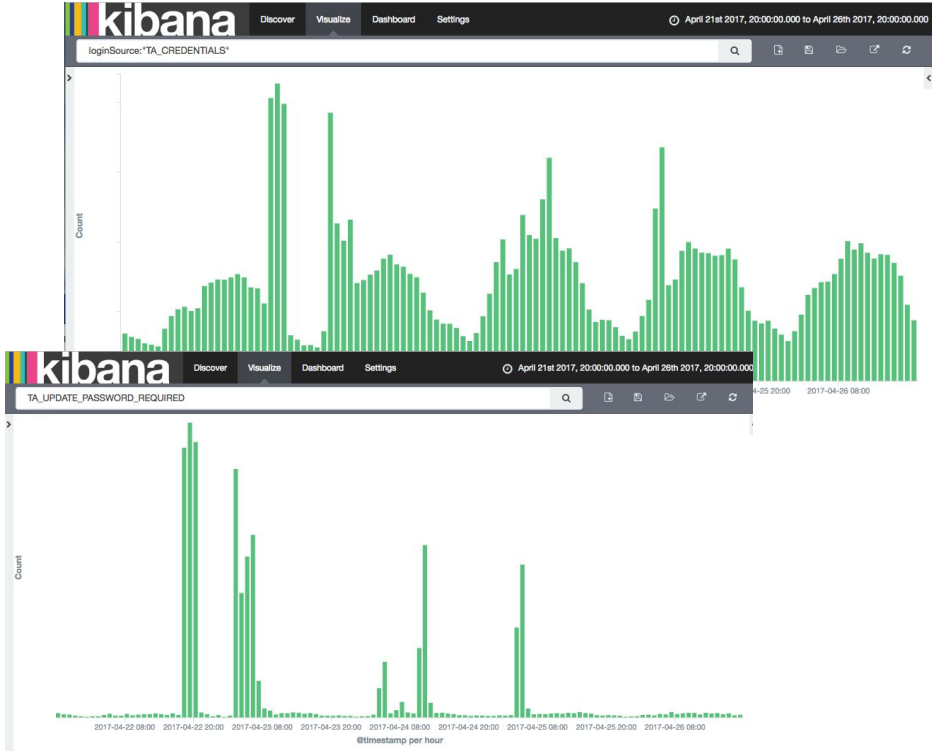

tripadvisor | Anatomy of a Probing Attack | Detection

- Kibana was invaluable for rapid data mining and investigation
  - Community-team Kibana - Ops dashboard and investigation tool
  - “ELK” stack: ElasticSearch, Logstash, Kibana
- The TA_UPDATE_PASSWORD_REQUIRED event saved the day, but wasn’t the best event to use as a signal.
- We added events for failed password attempts per e-mail and invalid e-mail login attempts (no such user).
  - Clear indicators of illegitimate behavior in bulk
  - These events were deliberately omitted in initial build; thought to be extraneous data. We hadn’t considered this scenario.

Anatomy of a Probing Attack | Detection

Now we have a consistent, clear signal, with alerting

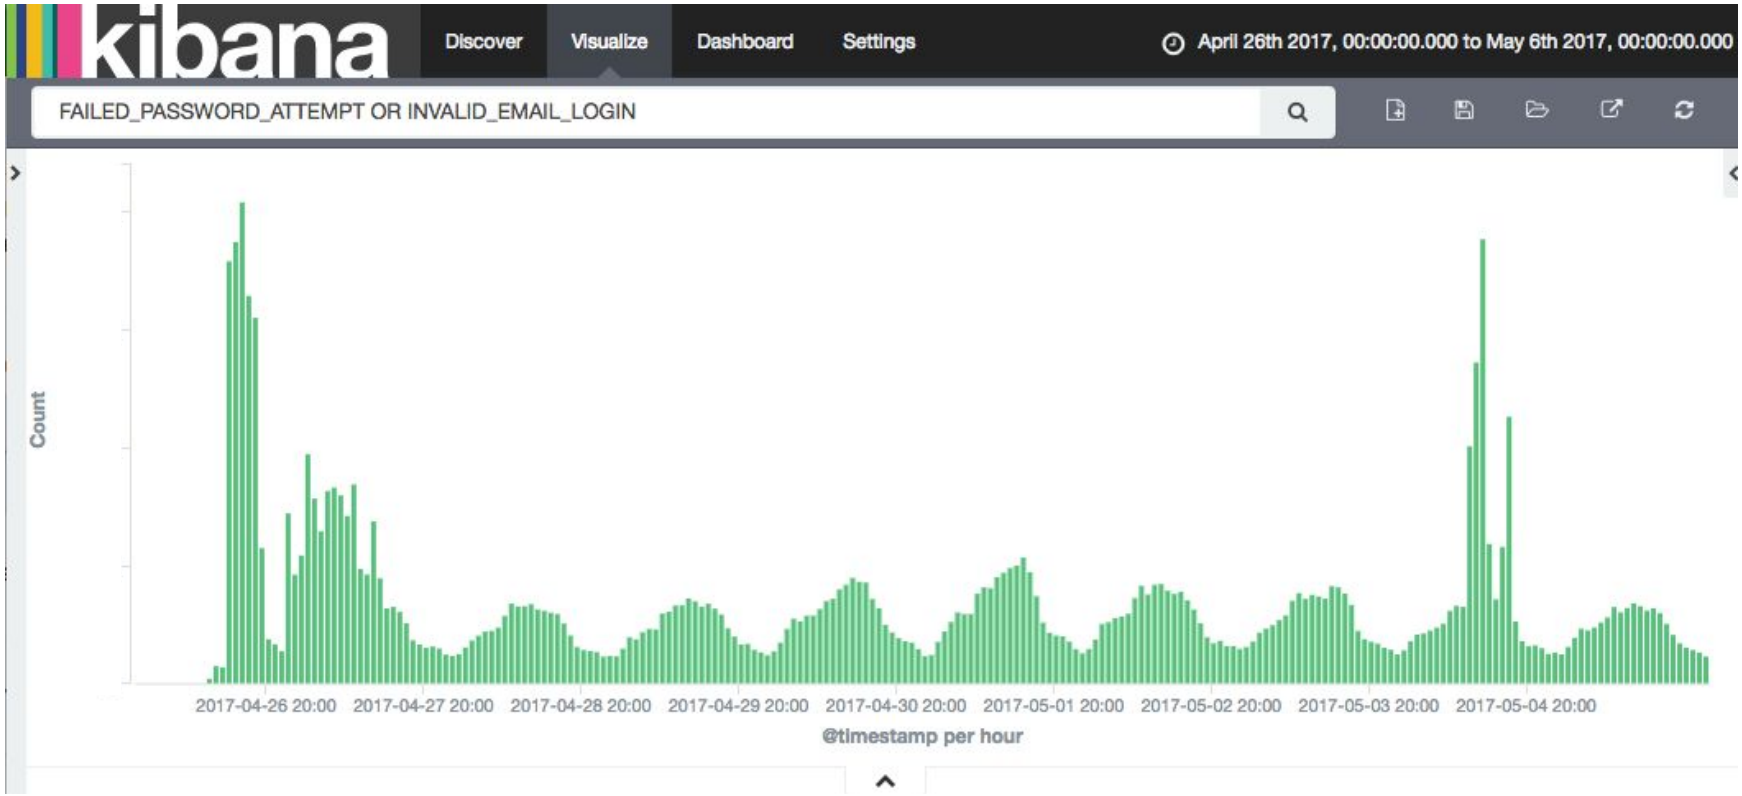

Anatomy of a Probing Attack | Thwarting

- Trusov attack
  - Identified the common network provider owner using “whois” on IPs
  - Forced Captcha on for the Russian Federation and Ukraine (53% of IPs)
  - Subnets were added to DoS filters, supplementing individual IPs
  - Ops implemented better DoS capabilities that allowed for CIDR* rules going forward
- User Agent (UA) attack
  - Added UA rewrite rules to block specific, illegitimate user agents
- These are small hurdles to overcome for a dedicated attacker
- Recent enhancements
  - Additional technical requirements for API registration or login
  - Invisible Captcha, always on for registration and login

* CIDR = Classless Inter-Domain Routing

Anatomy of a Probing Attack | Thwarting

Following the steps noted the previous slide, Trusov probed us again on 5/4. We improved blocking verification of pwned accounts by orders of magnitude.

tripadvisor® **Anatomy of a Probing Attack** | Remediation

- Identified compromised accounts by mining logs for successful logins using irregular patterns that included:
  - UserAgent string
  - IPs and subnets (identified by network provider ownership)
  - Omission or repetition of spoofed device id(s)
- All accounts were re-secured by resetting the password to a strong, random value, and invalidating logged-in sessions.
- Users were notified via e-mail of suspicious activity, and encouraged to change their passwords where they might be used elsewhere.
  - Some states and countries require this communication, though we all agreed it was the right thing to do.
  - The passwords were invariably used elsewhere, since there was no data leak at TripAdvisor

Analysis

- Logins issued were legit email/password credentials
  - Our brute-force attack security measures weren't tripped
  - Our network and server security wasn't compromised
- We determined passwords were not guessed
  - Used internal anti-fraud tools to determine there was a low frequency of similar passwords across the millions of requests
  - That conclusion was reinforced using the new FAILED_PASSWORD_ATTEMPT event to determine only 1-2 attempts were made per e-mail account
- There was no data breach
  - Our internal member DB contains many testing accounts that don't exist in the wild
  - These accounts can be used as "honeypots" to detect an internal breach. They were not probed.

tripadvisor

Analysis

- We were lucky
  - We detected the attack, because old accounts needed to update non-SSL-created passwords. The botnet caused those events to fire at an abnormal rate.
  - We historically had Logstash import errors that caused all ElastAlerts to fire with regularity. This made it difficult to tell a true signal from noise. We fixed this right around the time of the April probing.
- Probing had been going on periodically since November, 2016. April represented a huge ramp-up in volume.

Industry Security Notes

We met with the security team from another online travel website to trade notes. Our collective conclusions were:

- Captcha is the best protection against botnet attacks
- A WAF appliance (web application firewall) can protect your site from hackers and would-be criminals.
  - You need to train the appliance to recognize legit traffic patterns from irregularities.
  - It’s huge time investment to set up, requiring around a year of training to become reliable
  - The WAF automatically blocks traffic by IP/Subnet/CIDR when it detects abnormalities
- Probing attacks have ramped up significantly since November, 2016
  - Dealing with these events is the new “normal”

Take-aways

- You will be attacked
- Captcha is an easy-to-deploy security mechanism
- You ***must*** be able to measure and detect anomalies
- Visualizing anomalies via an Ops dashboard is super-helpful
- Automated alerts help you respond in near-real time
- Consider how your data is secured and gated.
  - If an account is compromised, don't leak critical information
  - Obfuscate email addresses, phone numbers, and credit card numbers
    - b……@tripadvisor.com
    - xxxxxxxxxxx5007, expires 04/20 (just like a receipt)

tripadvisor

Q&A

Questions?


tripadvisor®