Paper Review

Tony Espinoza

am.espinoza@utexas.edu

An Analysis of China's "Great Cannon"

- Background
- Great Cannon
- Evaluation
- History of Use
- Attribution the Great Cannon
- Potential Enhancements
- Conclusion

Background

- What is the Great Firewall (GFW,GFC)
  - On path system.
  - **NOT** in path.
  - What is the difference?
  - Can inject additional packets, they cannot prevent in-flight packets from reacing thier destination.
- How is it different than the Great Cannon?

Background

- What happens if a RST is injected?
- Would the end host still receive the packet sent from the host?
- How is packet inspection done?
  - GFW runs packet re-assembly and censorship logic in multipole parallel processes.
  - All packets on a connection go to the same process.

Great Cannon

- An in-path system
  - What can an in-path system do?
  - Can it suppress traffic?

Great Cannon

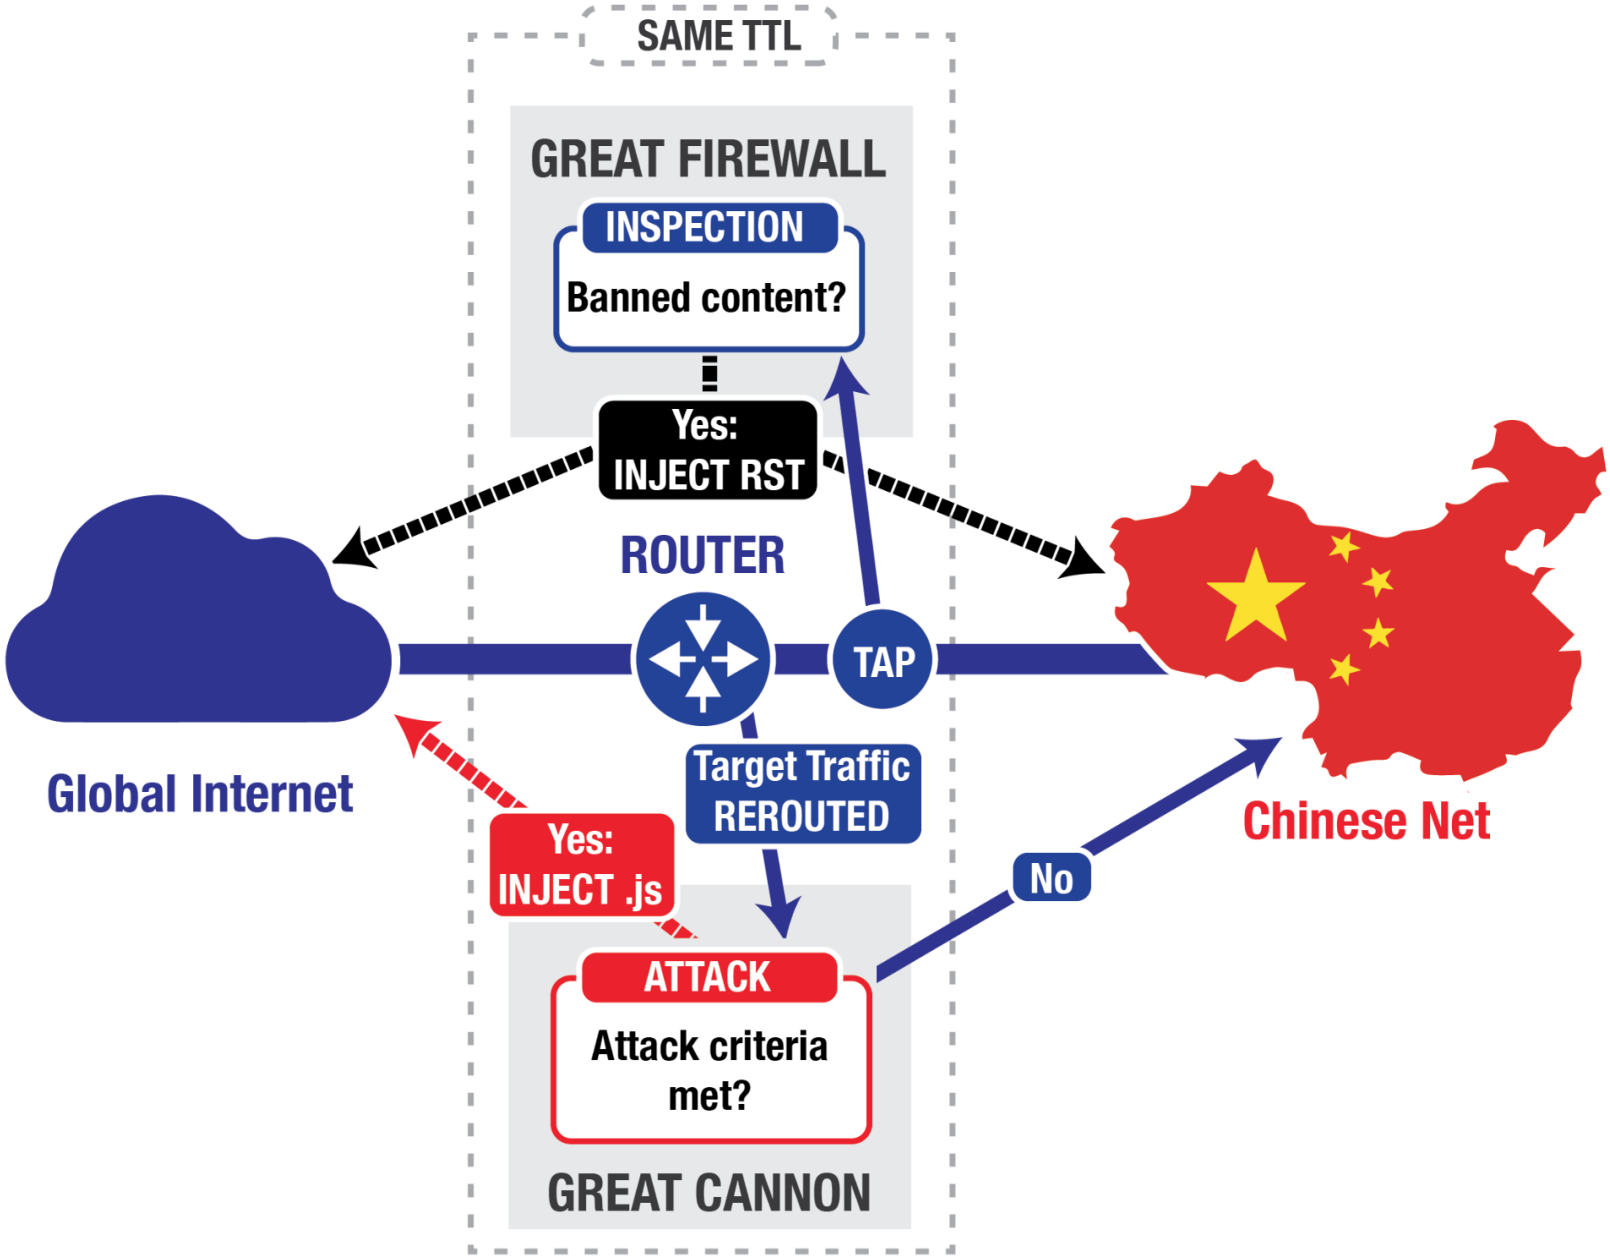

TTL Significance

Let's look at a pcap from homework 1.

Note:

- Censorship is dynamic
  - What works one day might not the next
  - http://www.baiwanzhan.com/
  - 法轮

Great Cannon

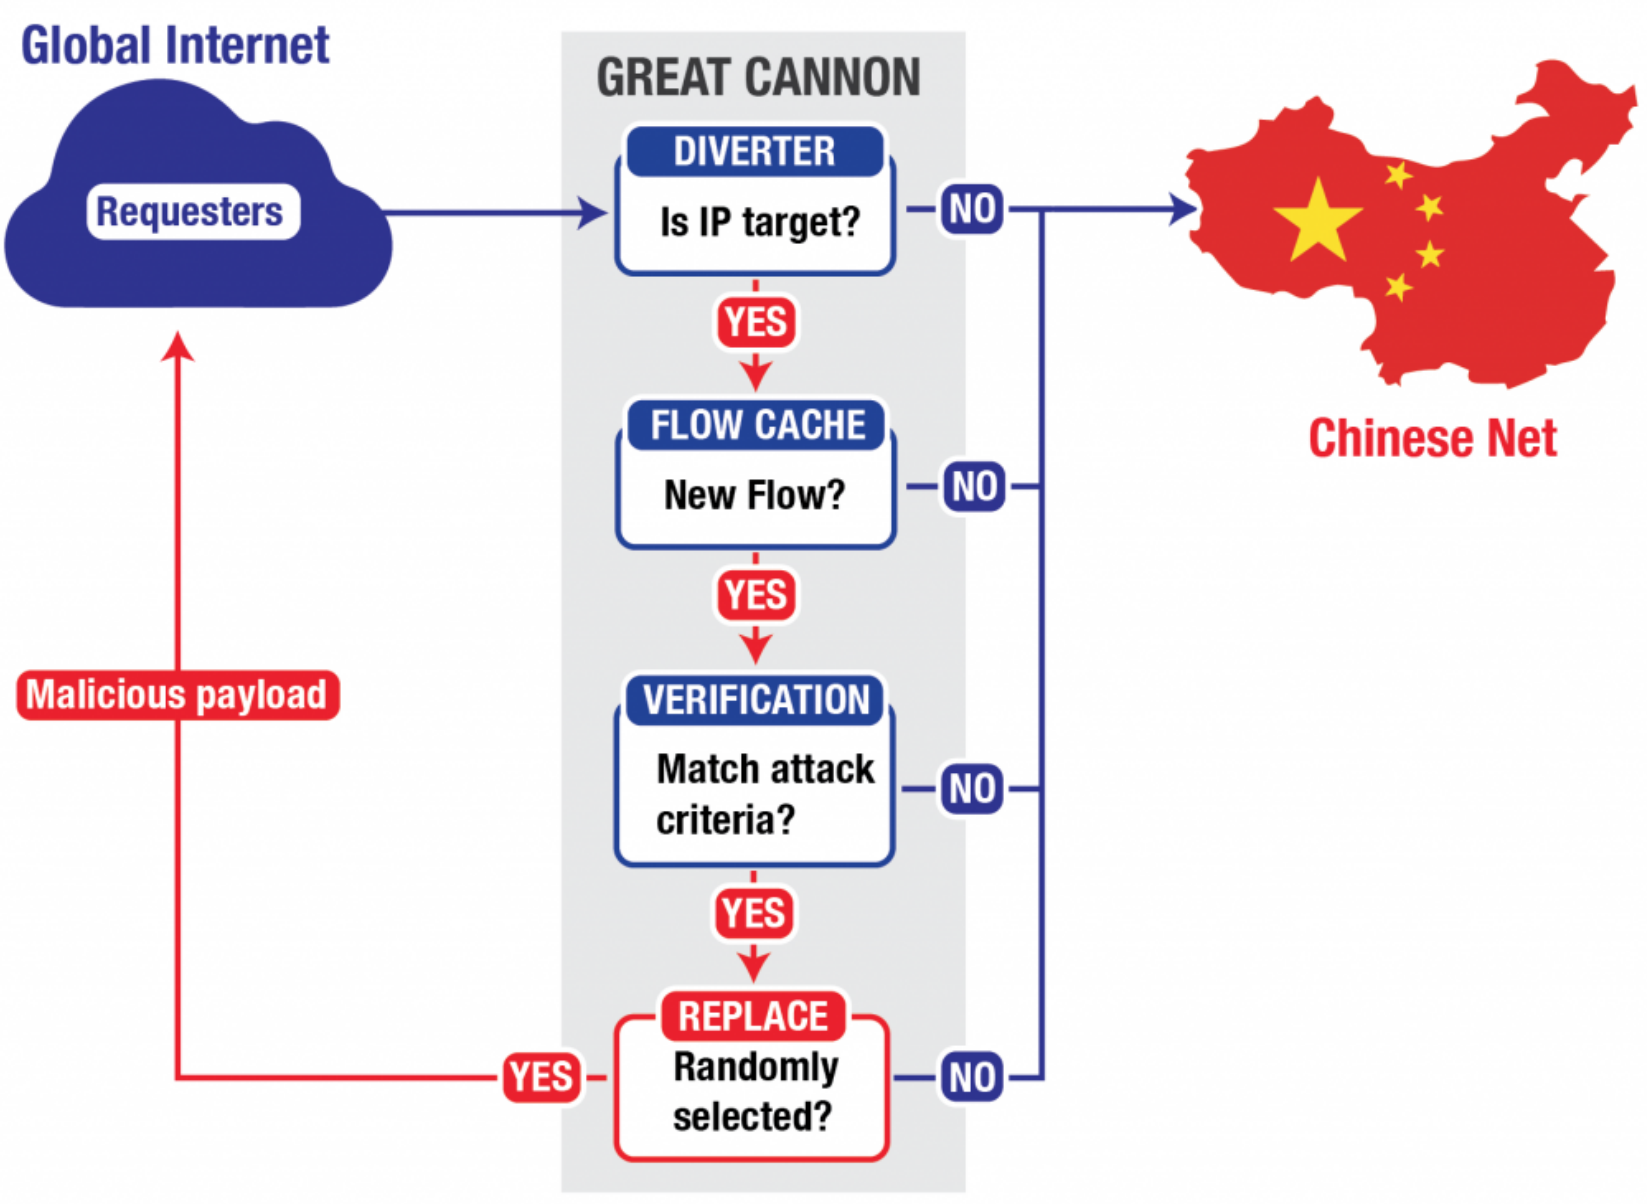

Great Cannon

- How was the in-path system used to perform a DDOS?
- Sent 1.75% of visitors to Baidu infrastructure services malicious javaScript.
  - Why only 1.75%?

Evaluation of Functionality

How did they verify that the GC:

- Appears to act probabilistically
- Operates as a separate in-path system
- Appears to be co-located with the GFW
- Was ‘aimed’ only at specific destination IP addresses

Appears to act probabilistically

- Tested from 4 different IP addresses.
- One IP was ignored by the GC.
- The other three the GC responded to 526 out of 30,000 requests.

Operates in-path

- GFW showed both the TCP RST as well as the legitimate server reply.
- GC does not show the server reply, only the injected malicious reply.

Co-located with GFW

- Used TTL to see where the GFW was located.
- Used TTL to see where the GC was located.
- How would you test this?

Was ‘aimed’ at specific destination IPs

- Tested on IP close to Baidu server.
  - GC ignored the request.
  - GFW acted on the censorable requests.

History of Use

- Google's Safe Browsing project captured instance of the attack.
  - March 3^{rd} - April 7^{th}.
- They build a pcap analyzer.
  - Ran on LBNL network data.
- Anonymous colleague shared data.
  - One year earlier.

Analysis of GreatFire.org Logs

- March 18^{th} 11:00 - March 19^{th} 7:00 GMT.
- Each hour randomly select ~30MB of compressed logs.
- Use MaxMind GeoIP2 Lite DB.
  - MaxMind is a IP address to lat,long database.

Analysis of GreatFire.org Logs

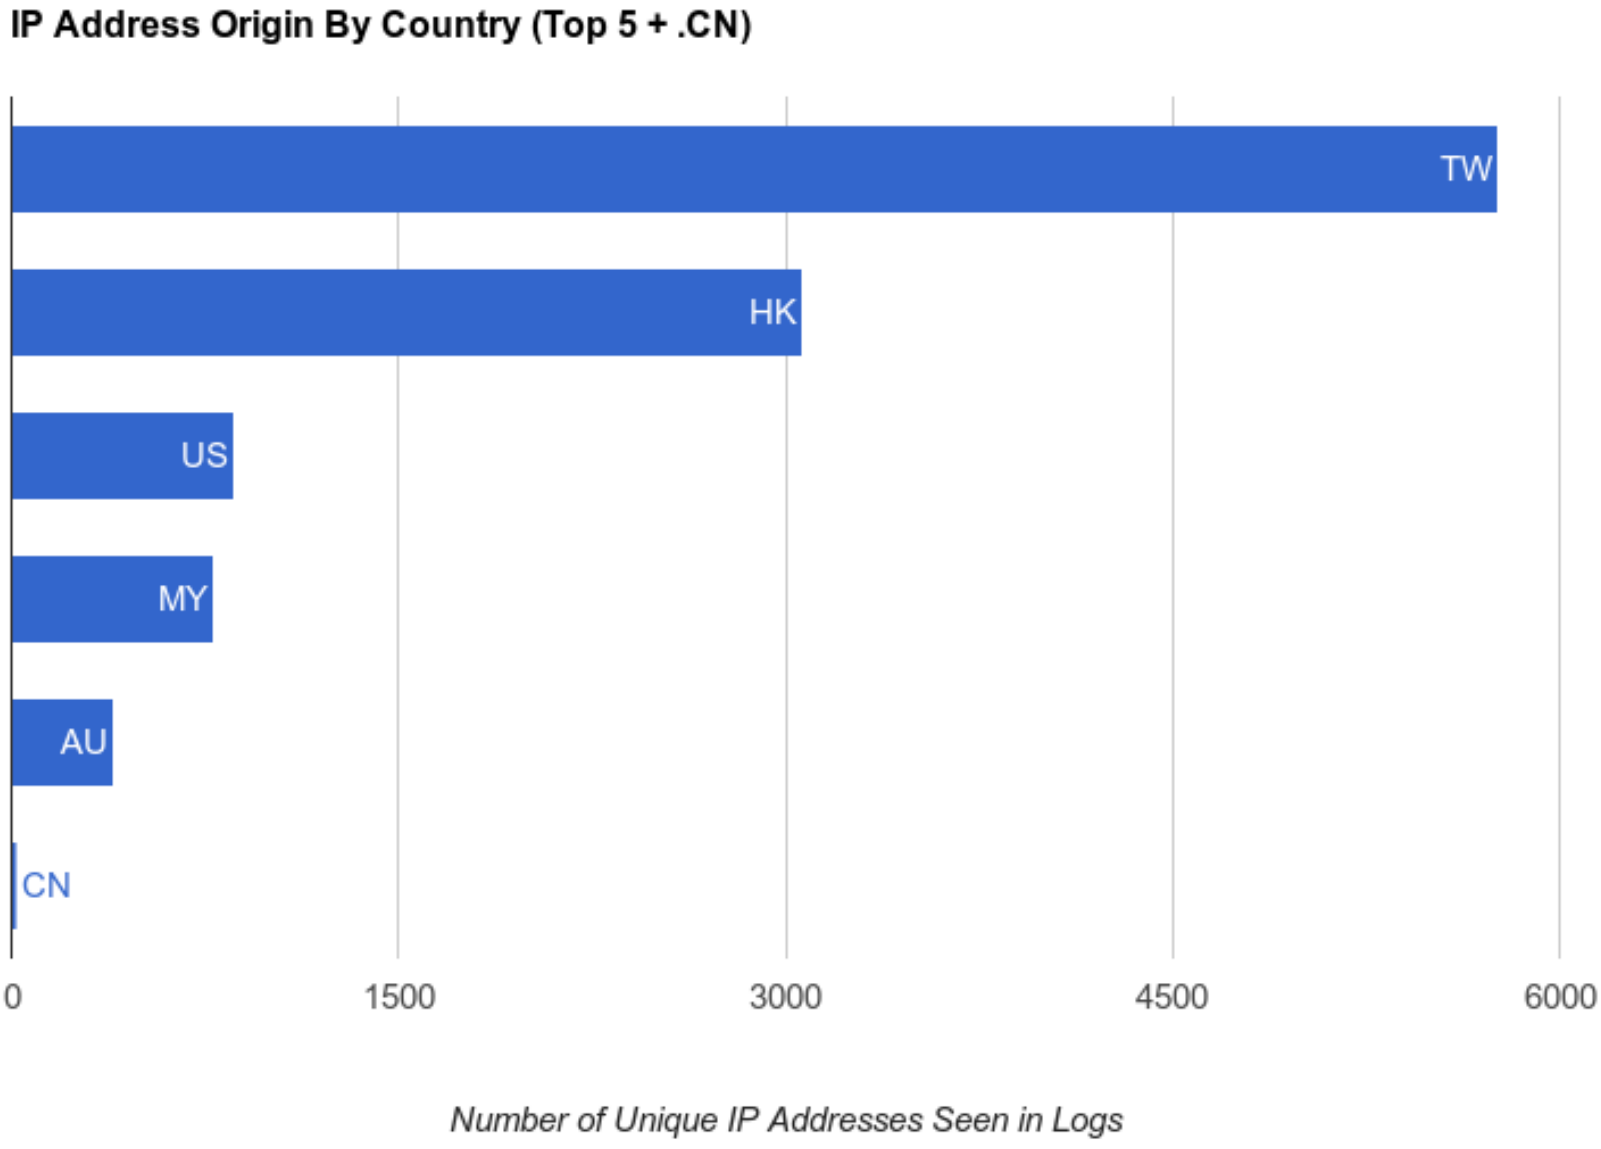

Analysis of GreatFire.org Logs

- Why are TW and HK so high up?
- Why does CN not have more unique IPs?

Attribution the Great Cannon

- Where does the GC operate?
- Who built the GC?
- What is its use?
- Who was it attacking?

Where does the GC operate?

- co-located with the GFC.
  - Tested from 2 different international Internet links into different Chinese ISPs.
    - Found the GFW and GC were co-located in both.
  - Suggests a governmental actor.

Who built it?

- The authors suggest the same architect of the GFW.
  - Both have similar behavior of TTLs.

What is its use?

- Not suited for censorship.
  - Why?
  - Only looks at first packet.
  - Targets specific IP destinations.
- MITM to inject traffic.

Who was it attacking?

- GreatFire.org
  - Service targeted provides proxies to bypass the GFW (CloudFront)
- GitHub
  - Hosted 2 GreatFire.org repos
  - Why attack github and not block it?
    - They tried to, but got negative reaction.

Potential Enhancements

- Can switch targeting *source* IP, to *destination* IP and target individuals.
- Could fix its network artifacts.
  - Make it harder to detect.
- Could be used to intercept email.

Questions

- Would this work with HTTPS connections?

What next?

- In the CG paper the infrastructure is set up to aid a nation state.
- What can be done when there is no infrastructure?
- Can we exploit the protocols?
  - Yes.
  - How?

Exploiting protocols

- Can we scan a machine without giving ourselves away as the scanner?
- Can we find a machine behind a firewall?
- We can see if two machines are communicating.
  - Trivial if you are on the path between the two hots.
- What about if you are off path?
  - What would off path measurement look like?

Network side channels

- What is a side channel?
- What do I mean by network side channel?
  - A side channel in the implementation of a networking protocol.

IPID side channel

- One way to choose an IPID is globally incrementing
  - What does this look like?
- What can we do with a globally incrementing IPID?
  - Idle scan.

Idle scan

- A scan technique where you don't use your IP to scan the victim.
- Setup
  - Zombie
  - Victim
  - Attacker
- Proposed by Antirez in 1998

Idle scan

- Caveats
  - Not many globally incrementing IPIDs any more.
  - Machine must be idle so you don't have noise.
- Overcome caveats?

Scan through firewall?

- Use SYN cache as a shared limited resource
- Fill the SYN cache
- Send spoffed SYN packet, two outcomes.
  - Get RST
    - Room is SYN chache.
    - Attacker will get a SYN/ACK.
  - Get nothing
    - SYN cache is full
    - Attacker will get a SYN cookie.