




Weapons of mass intelligence to fight the bad bots


**Stefan Mardak,
Senior Enterprise Security Architect**

Your shop in the internet - The real world



Leveraging Big Data to Understand Attackers


The following slides are based on a real events on January 5th 2014....



“Akamai, we are under attack!...”

Ad-Hoc Attack Analysis

An attempt to exploit an old (2007) WordPress Remote File Inclusion vulnerability. The victim application was running ASP.NET



```
GET /wp-content/wordtube-button.php?wpPATH=http://www.google.com/humans.txt?  
HTTP/1.1  
Host: www.vulnerable.site  
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_8_4)
```

Attacked parameter : wpPATH

Malicious payload: http://www.google.com/humans.txt

What Else Did This Attacker Do On This Site?

Same attacker sent **2122** different RFI exploit attempts



34 different sites were attacked by the same attacker
with a total of 24,301 attacks



Was There Similar Activity Going On At The Same Time?

Attacks originated from a **botnet** containing **272** attacking machines

1696 victim applications were targeted

1,358,980 attacks were launched during the campaign

The campaign lasted for **2** weeks

Taking Big Data to the Next Step...

Forecast malicious intent *before* exploitation

Take action on malicious clients

Integrate with other intelligence systems



The Akamai Intelligent Platform



The Platform

- 175,000+ Servers
- 2,700+ Locations
- 750+ Cities
- 108 Countries
- 7 Continents
(including Antarctica)
- 1,300+ Networks

The Data

- 2 trillion hits per day
- 780 million unique IPv4 addresses seen quarterly
- 13+ trillion log lines per day
- 260+ terabytes of compressed daily logs

15 - 30% of all web traffic


Record past behavior — use data to protect everyone

- Analyze activity over Akamai customers
 - Delivery & Security
- Identify bad clients based on past behavior
- Attach a risk score to bad clients
- Take action based on risk score




CSI Architecture

Hadoop Cluster



CSI Architecture




1. WAF triggers
2. CDN logs – “WAF light”
3. CDN logs – behavioral

The challenge

- Log all security events > 2 billion WAF triggers/day
- 13 trillion log lines a day....how to parse through that?

CSI Architecture

Data Sources



1. WAF triggers
2. CDN logs – “WAF light”
3. CDN logs – behavioral

Hadoop Cluster



Heuristics



1. Attack patterns
2. Client behavior
3. Application profiling
4. Shared IP detection
5. False positive reduction



CSI Platform Statistics

2 Petabytes of security data stored

20 Terabytes of daily attack data

Retention for up to 90 days

Client Reputation Categories & Scores

Categories

- Web attacker: participating in web vulnerability exploitation attempts
- DoS attackers: participating in DoS attempts
- Scanning Tools: performing vulnerability scanning
- Web Scrapers: extract information from the application

Risk score ranges from 1-10, with 10 being high, based on:

- Persistency of the attacker
- Severity of the attack
- Magnitude of the attack
- Distribution of the attack across multiple hosts




Client Reputation Details

118.103.239.5

Load

Add to Whitelist

REPUTATION OVER TIME - 05/26/2014, 07:35:36 PM - 05/28/2014, 07:35:36 PM



Web Attackers

SCORE CHANGING EVENTS

Refresh

TIME	CATEGORY	BEFORE	AFTER	REASONING
05/28/2014 - 08:19:00 AM	Web Attackers	10	9	Client risk score decay
05/27/2014 - 05:20:00 PM	Web Attackers	9	10	Client performed 1549 SQL injection attempts using 37 unique attack payloads
05/27/2014 - 11:19:00 AM	Web Attackers	8	9	Client performed 691 SQL injection attempts using 18 unique attack payloads
05/27/2014 - 10:22:00 AM	Web Attackers	4	8	Client performed 232 SQL injection attempts using 9 unique attack payloads
05/27/2014 - 06:19:00 AM	Web Attackers	5	4	Client risk score decay

Client Reputation Details

118.103.239.5

Load

Add to Whitelist

REPUTATION OVER TIME - 05/26/2014, 07:35:36 PM - 05/28/2014, 07:35:36 PM

2d 4d 6d 8d All

5

0

27. May

08:00

09:00

10:00

10:30

10:45

11:00

11:30

17:30

28. May

30. May

21. May

25. May

27. May

11:00

18:00

30. May

31. May

3. Jun

Web Attackers

SCORE CHANGING EVENTS

Refresh

Risk score decay

		BEFORE	AFTER	REASONING
		10	9	Client risk score decay
		9	10	Client performed 1549 SQL injection attempts using 37 unique attack payloads
		8	9	Client performed 691 SQL injection attempts using 18 unique attack payloads
05/27/2014 - 10:22:00 AM	Web Attackers	4	8	Client performed 232 SQL injection attempts using 9 unique attack payloads
05/27/2014 - 06:19:00 AM	Web Attackers	5	4	Client risk score decay

Client Reputation Details

118.103.239.5

Load

Add to Whitelist

REPUTATION OVER TIME - 05/26/2014, 07:35:36 PM - 05/28/2014, 07:35:36 PM

2d 4d 6d 8d All

5

0

27. May

08:00

09:00

10:00

10:30

10:45

11:00

11:30

17:30

28. May

30. May

4

21. May

25. May

27. May

11:00

18:00

30. May

31. May

3. Jun

5

Web Attackers

SCORE CHANGING EVENTS

Refresh

TIME

05/28

05/27

05/27

05/27/2014 - 10:22:00 AM

Web Attackers

4

8

Client performed 232 SQL injection attempts using 9 unique attack payloads

05/27/2014 - 06:19:00 AM

Web Attackers

5

4

Client risk score decay

1549 SQL injection attempts w/37 unique payloads

Client Reputation Details

118.103.239.5

Load

Add to Whitelist

REPUTATION OVER TIME - 05/26/2014, 07:35:36 PM - 05/28/2014, 07:35:36 PM

2d 4d 6d 8d All

5

0

27. May

08:00

09:00

10:00

10:30

10:45

11:00

11:30

17:30

28. May

30. May

21. May

25. May

27. May

11:00

18:00

30. May

31. May

3. Jun

4

Web At

SCORE CHANGING EVENTS

Refresh

TIME	CATEGORY			
05/28/2014 - 08:19:00 AM	Web Attackers			
05/27/2014 - 05:20:00 PM	Web Attackers			
05/27/2014 - 11:19:00 AM	Web Attackers			
05/27/2014 - 10:22:00 AM	Web Attackers	4	8	Client performed 232 SQL injection attempts using 9 unique attack payloads
05/27/2014 - 06:19:00 AM	Web Attackers	5	4	Client risk score decay

691 SQL injection attempts w/18 unique payloads

Client Reputation Details

118.103.239.5

Load

Add to Whitelist

REPUTATION OVER TIME - 05/26/2014, 07:35:36 PM - 05/28/2014, 07:35:36 PM

2d 4d 6d 8d All

5

0

27. May

08:00

09:00

10:00

10:30

10:45

11:00

11:30

17:30

28. May

30. May

21. May

25. May

27. May

11:00

18:00

30. May

31. May

3. Jun

Web Attackers

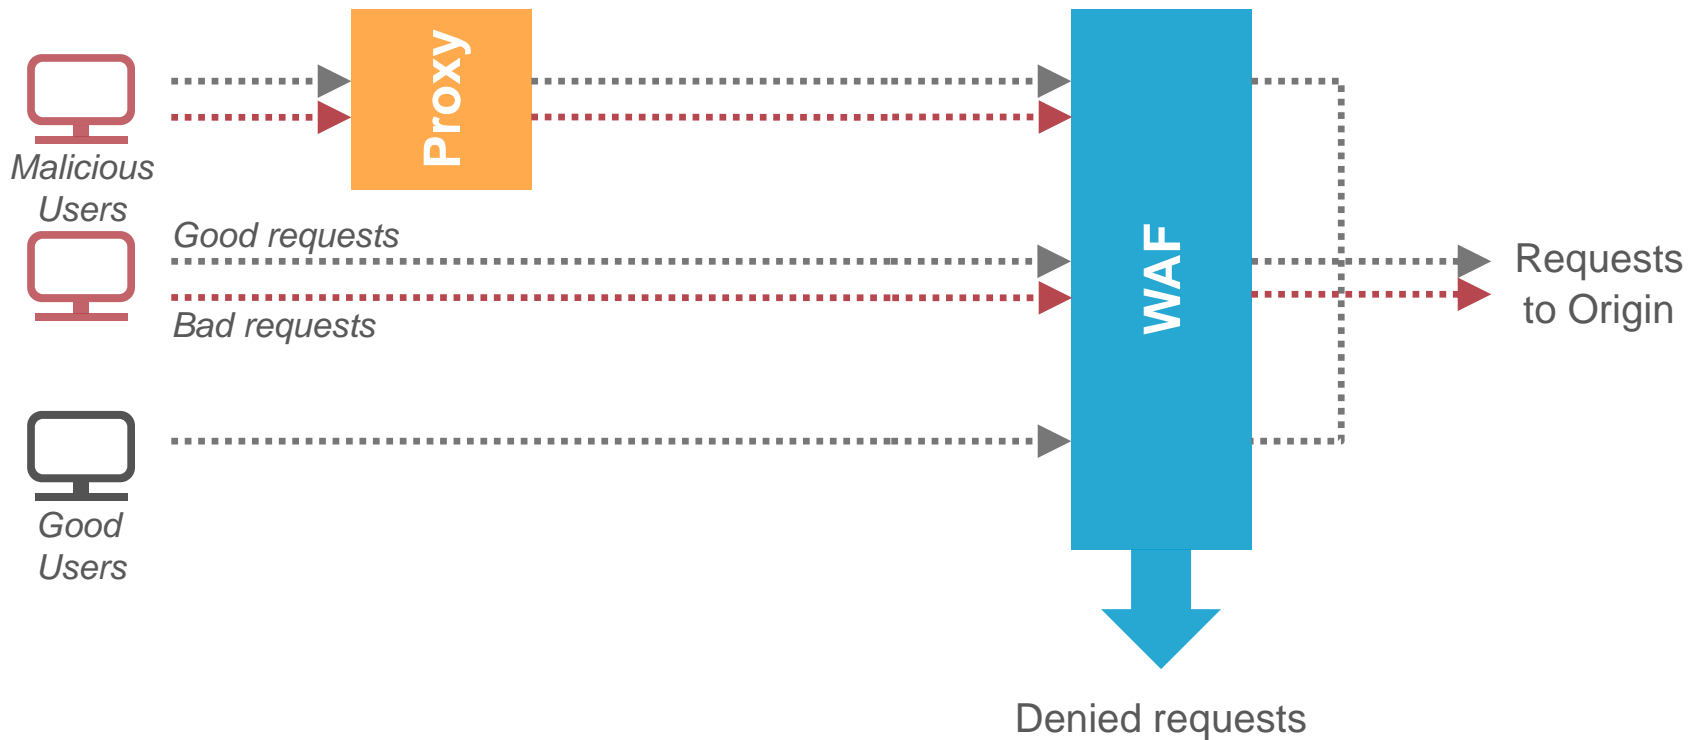
SCORE CHANGING EVENTS

Refresh

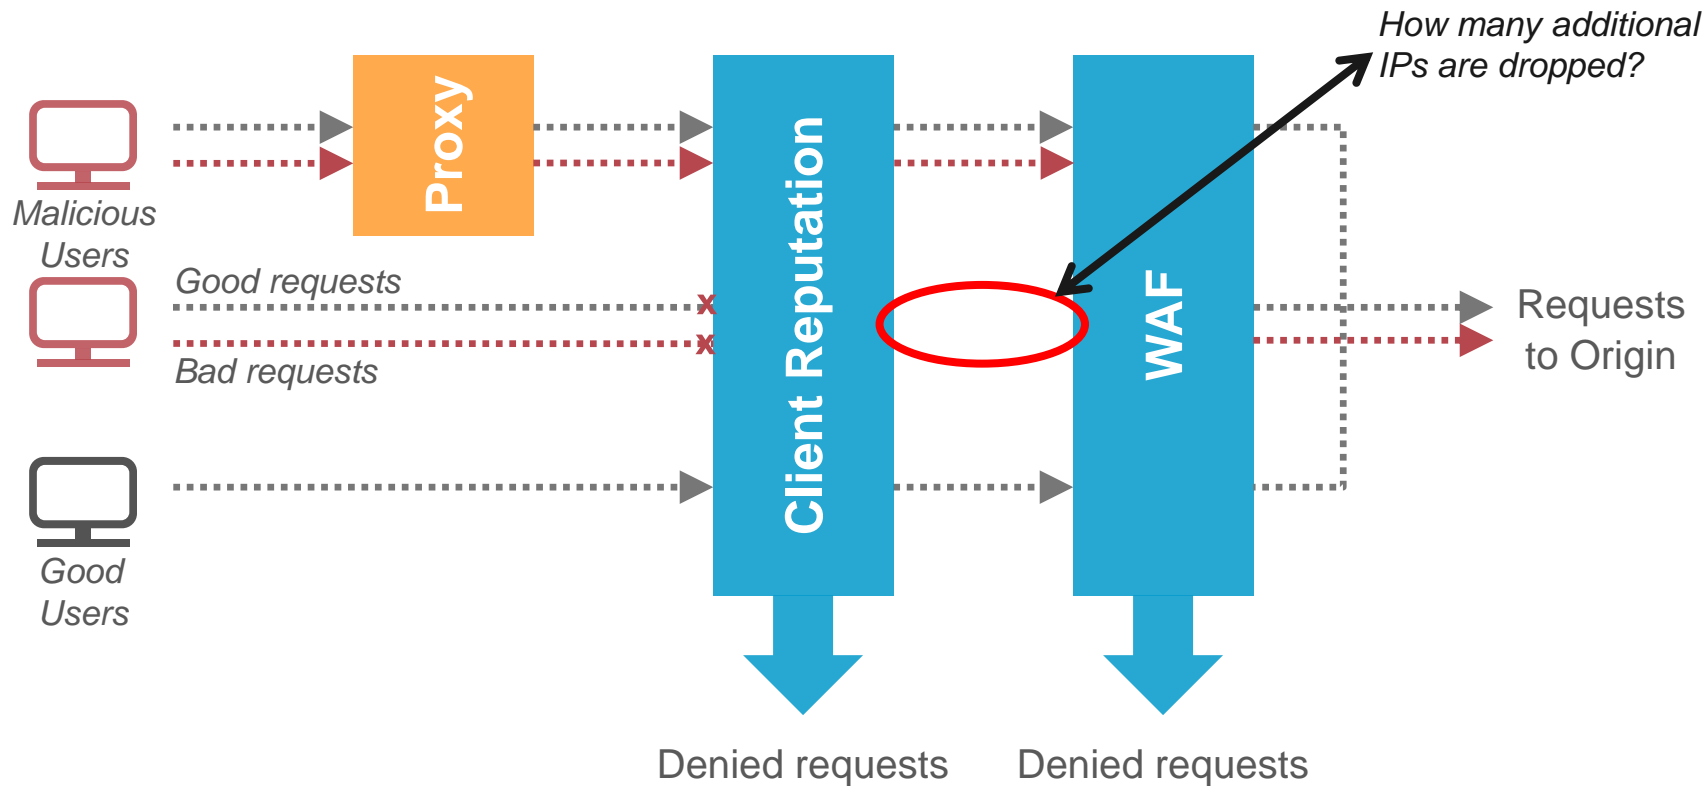
TIME	CATEGORY	BEFORE	AFTER	REASONING
05/28/2014 - 08:19:00 AM	Web Attackers	10	9	Client risk score decay
05/27/2014 - 05:20:00 PM	Web Attackers	9	10	Client performed 15
05/27/2014 - 11:19:00 AM	Web Attackers	8	9	Client performed 68
05/27/2014 - 10:22:00 AM	Web Attackers	4	8	Client performed 232 SQL injection attempts using 9 unique attack payloads
05/27/2014 - 06:19:00 AM	Web Attackers	5	4	Client risk score decay

232 SQL injection attempts w/9 unique payloads

Client Reputation Compliments WAF



Client Reputation Compliments WAF



Client Reputations Added Value

(“Why didn’t WAF catch this?”)

- CR sees ALL activities that end-users do (even the “good” ones)
- CR correlates events across time & “space” (sources & targets)
 - Distributed attack sources can easily get picked up as a “botnet”
 - Slow & low brute force attacks can be spotted over longer time periods
 - Detects scanning / searches for potentially vulnerable files (even 0-days!)
- CR can “connect the dots”
 - Many suspicious activities grouped together become a strong enough signal

-> Based on Client Reputation malicious users
can be blocked BEFORE the reconnaissance phase

Summary

Client Reputation is the Internet's criminal track record

Client Reputation focuses on the source of the attack

Forecasts the likelihood of a client to participate in an attack

Client Reputation provides insights to activities that sometimes elude WAF's transactional view

Client Reputation complements traditional detection techniques



Reference: <http://www.StateOfTheInternet.com>

Stefan Mardak – Senior Enterprise Security Architect
SMardak@Akamai.com


Further reading...

- Internet of things - how many traffic originates from bots (good or bad)?
- Attackers are using Bots for attacks – what do we know about botnets?
- Is a botnet local, regional or global?
- Attacking bots should be reported and then cleaned up – does this happen?
- Are cleaned up members replaced?
- How does attack traffic change over the course of time?

Bots on the Akamai Platform

8.01 BILLION

Bot requests in 24-hours

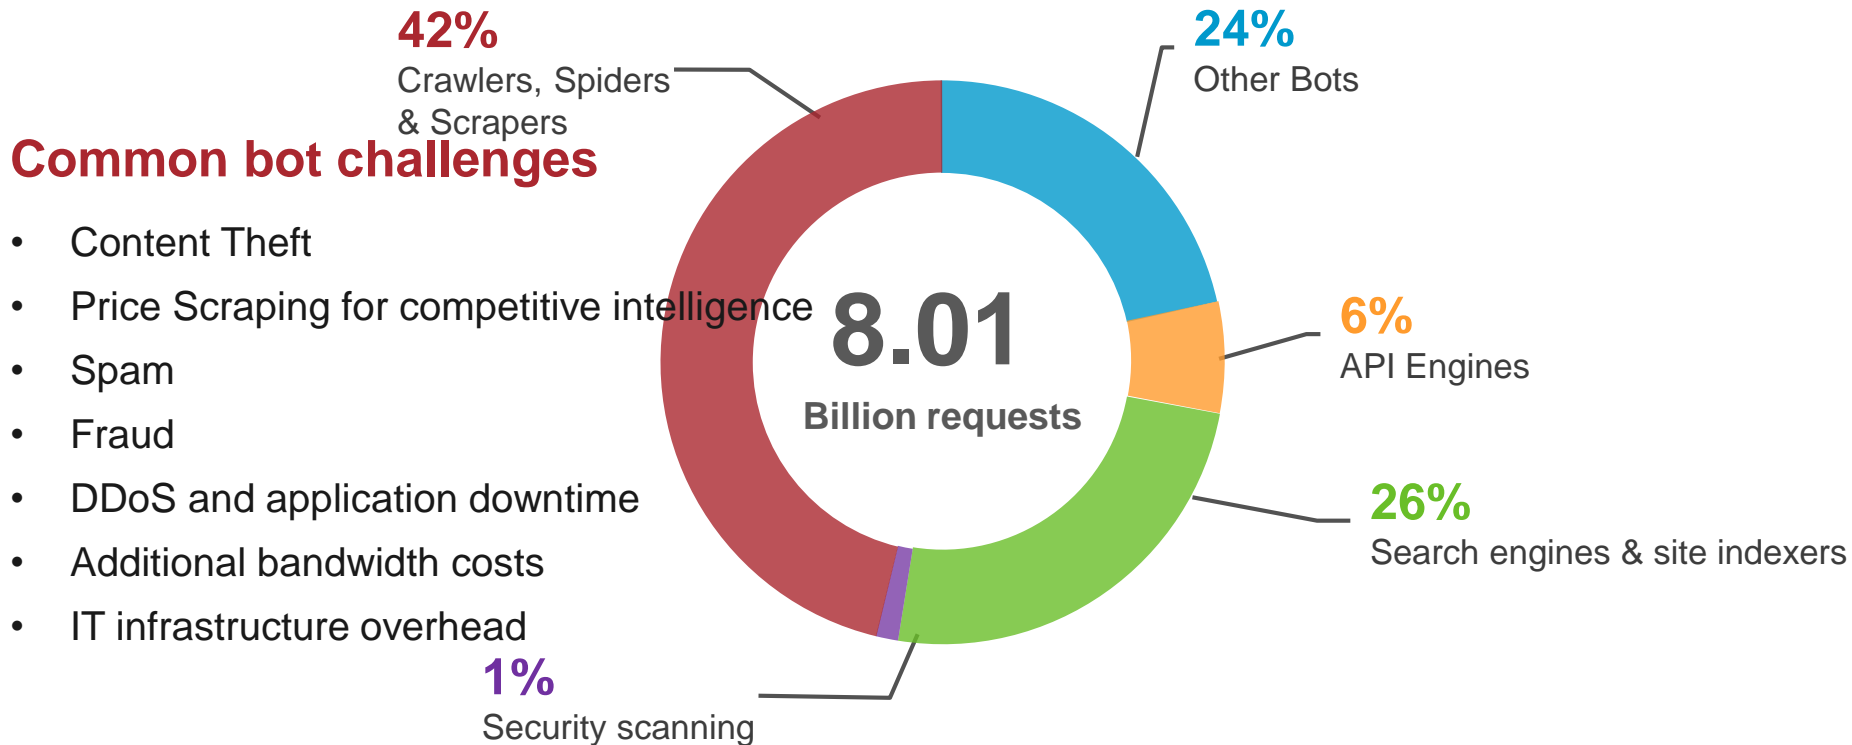


Data Collected
April 1-2, 2015

Total Requests:
85,475,034,620

Bots were 9.4% of all requests

Bots – The Akamai Viewpoint



A Year in the Life of a Botnet

In January 2014 we published a blog on a global botnet:

- <https://blogs.akamai.com/2014/01/analyzing-a-malicious-botnet-attack-campaign-through-the-security-big-data-prism.html>


Exploiting Joomla Content Editor vulnerability to install backdoors

Began as a “single event” analysis of the exploit


“Zoomed out” and discovered an entire botnet mining the web for vulnerable Joomla servers

A Truly Global Botnet

Botnet Machine Distribution by Country (Top 10)




Botnet Machine Distribution by Continent




And a Very Active Botnet

- 43,000 malicious HTTP requests seen over the month
- 2008 different web applications were targeted

Number of Attacks Per Day



Number of Targets Per Day



10 months later, the Botnet lives on...

In Nov. 2014, the team began a 3 month follow on analysis


The botnet now contains 1037 members.

All members are compromised public Web servers, mostly running Joomla and WordPress CMS


The Botnet has targeted more than 7800 applications over the period

Note – the data is only based on Akamai customers – probably targeted many more applications

Active Members Over Time



New Botnet Members Over Time



Activity Duration of Botnet Members and Evolution

On average, Joomla botnet members spurted malicious traffic over 29 days.

Comparing the active Botnet members from 9 months ago to now

- 43 of the botnet members were also maliciously active 9 months ago.
- 4% of botnet members have not been “cleaned up” for 9 months

The Botnet evolved over time to attempt to also exploit other vulnerabilities:

- Remote File Inclusion (RFI) on the TimThumb image resizer WordPress module
- Remote Code Execution (RCE) on the Open Flash Chart library