XINDI
BOTNET
pixalate
BOTNET REPORT

EXECUTIVE SUMMARY

Pixalate, the enterprise analytics and security platform, is issuing a cybersecurity threat advisory about the discovery of **Xindi botnet**, which will, if not stopped, cost advertisers an estimated **$3 billion by the end of 2016**.

Xindi, the first-of-its-kind botnet, was specifically developed by fraudsters to exploit a critical vulnerability (CVE-2015-7266) in the Internet advertising protocol implementation (OpenRTB) by turning enterprise and university networks into botnets that launch attacks on advertising exchanges. Xindi botnet, unlike its predecessors, does not defraud by clickjacking, which can be easily detected. Rather, it is the first botnet that **exclusively focuses on exclusively generating fake "viewable" impressions at scale**.

Pixalate's study reveals that Xindi Bot is a malware that has rapidly grown and spread over the past year, infecting millions of machines in operation at thousands of organizations. It appears that as many as **6-8 million machines** in more than 5000+ reputable organizations, including 10% of Fortune 500 organizations, 1500 university networks, and more than 200 financial and government organizations, are at risk, making Xindi one of the most significant enterprise-level botnets engaged in distributed advertising fraud attacks.

Hackers have for years been using various deceiving techniques to make money from advertising fraud. But Xindi marks a paradigm shift in how the world of digital advertising is now directly impacted by cyber-security breaches. Pixalate is releasing this advisory for Chief Information Security Officers (CISOs) of Enterprises and Universities with details on how Xindi Bot operates and measures they can take to restrict this botnet.

	XINDI	CHAMELEON	ZEROACCESS
DATE OF DISCOVERY	JULY 6TH, 2015	FEBRUARY 28TH, 2013	MAY 2011
ATTACK TYPE	IMPRESSION FRAUD	CLICK FRAUD	CLICK FRAUD
MEDIA	DESKTOP, MOBILE	DESKTOP	UNKNOWN
NUMBER OF MACHINES	6-8 MILLION	120,000	2.2 MILLION (2012)
HOST MACHINE TYPE	REPUTABLE ORGANIZATIONS, ENTERPRISES & UNIVERSITIES	RESIDENTIAL	UNKNOWN

COPYRIGHT © 2015 PIXALATE

ANATOMY OF THE BOTNET

COPYRIGHT © 2015 PIXALATE

HOW DOES XINDI BOTNET WORK?

Xindi is the first botnet of its kind to attack and exploit the OpenRTB protocol implemenation in order to gain profit from programmatic advertising. Unlike other botnets profiteering from click fraud, Xindi conducts impression fraud from inside high-value enterprise networks.

Xindi stakes out enterprise machines for financial gain and operates under the radar, accessing large amounts of bandwidth and the whole machine's computing power when the user is away. These networks are less susceptible to being blocked on the IP level because they generally have good reputations and far higher CPMs.

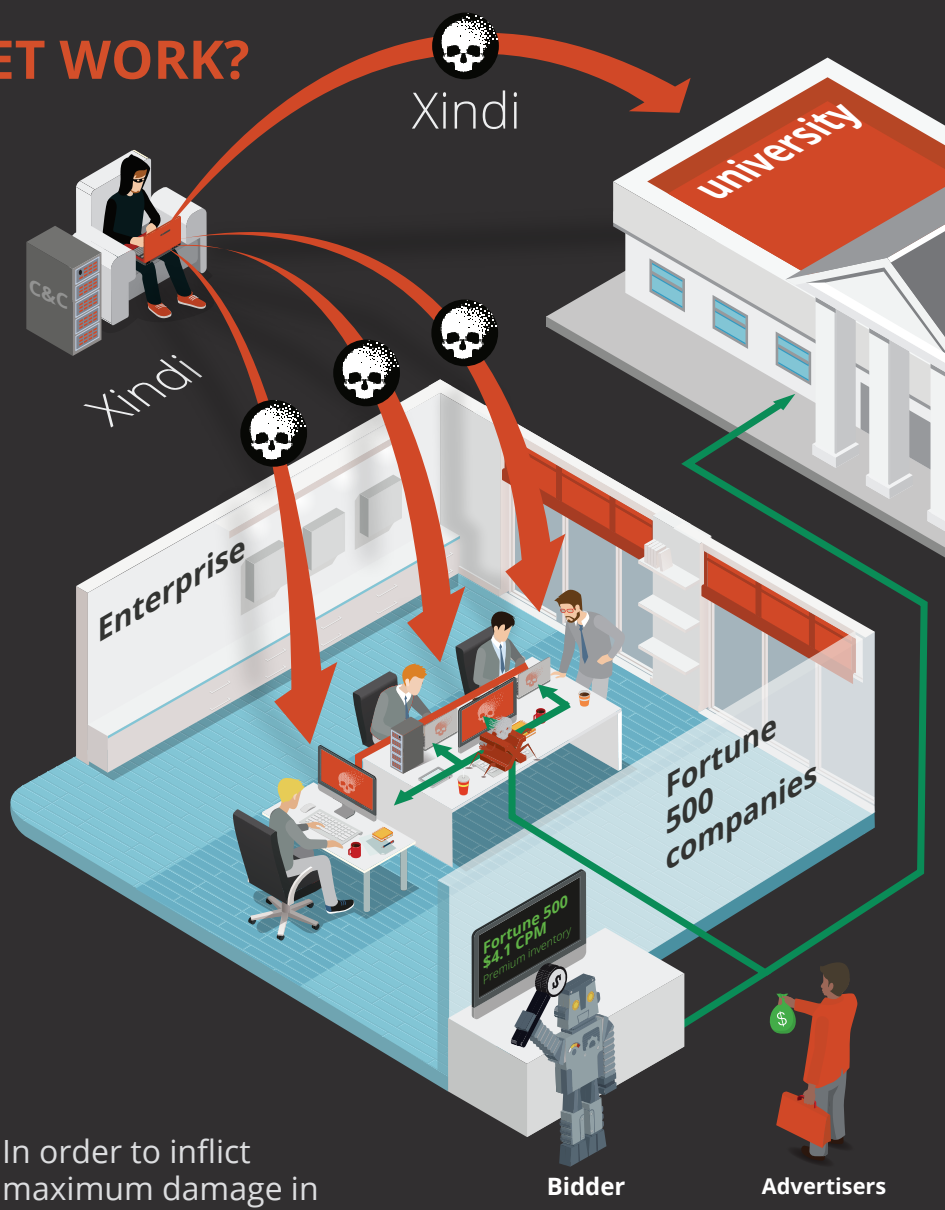

FIGURE 01

FIGURE 02

In order to inflict maximum damage in the shortest span of time, Xindi systematically exploits the Amnesia vulnerability. This vulnerability allows Xindi to conceal the true status of an ad transaction, which in turn causes bidding engines to bid on more impressions per compromised host than originally intended. Xindi achieves this by hoarding multiple ad markups in a transient state for hours on end and replaying them in a burst. In active campaigns the botnet was found to increase fraud by up to 300% and to consistently falsify viewability by up to 90%.

COPYRIGHT © 2015 PIXALATE

AMNESIA BUG

WHAT IS THE AMNESIA BUG?

The Amnesia Bug is a critical vulnerability (CVE-2015-7266) in the OpenRTB v2.3 protocol implementation, which is the standard for real-time digital media buying and selling. This vulnerability allows fraudsters to conceal the true status of an ad transaction, which in turn causes bidding engines to bid on more impressions per compromised host than originally intended. Fraudsters achieve this by hoarding multiple ad markups in a transient state for hours on end and replaying them in a burst. This has the potential to corrupt the bidding logic and compromise the integrity of the bids.

Pixalate has observed that as a result of fraudsters exploiting the Amnesia Bug, fraud on affected campaigns increases by up to 300 percent. The affected campaigns tend to have high viewability (in the range of 85-95 percent) and a highly desired user base. Frequency cap controls fail to work, and discrepancies between ad exchanges and demand-side platforms (DSPs) spike.

> "FRAUDSTERS HAVE EXPLOITED THE AMNESIA BUG, CAUSING FRAUD TO INCREASE BY UP TO 300 PERCENT ON AFFECTED CAMPAIGNS."

WHAT IS OPEN RTB?

The advertising ecosystem has radically evolved in the past few years in pursuit of efficiency and growth. This evolution has been made possible by the automation of ad buying and selling systems, i.e., programmatic advertising, which will account for 52% of non-search digital ad spend globally by the end of 2015, per eMarketer.

OpenRTB is an open communications standard for real-time machine-based buying of advertising that is used by leading companies.

COPYRIGHT © 2015 PIXALATE

HOW DOES OPEN RTB NORMALLY WORK?

When a user visits a publisher site that contains ad placements, the browser requests an ad from the ad-serving platform by passing it all relevant information about the user and placement. For each inbound ad request, the supply-side platform (SSP) or exchange broadcasts bid requests to multiple DSPs. Upon receiving the bid request, the bid engines evaluate its details and decide if they want to send a bid response back to the SSP/exchange. The responses are then evaluated under the prevailing auction rules on the SSP/exchange side.

Once the auction is won, the ad creative is sent back to the browser. Once the creative is rendered, the winning DSP is notified. In some cases the winning DSP is also notified as soon as the auction is closed on the SSP side.

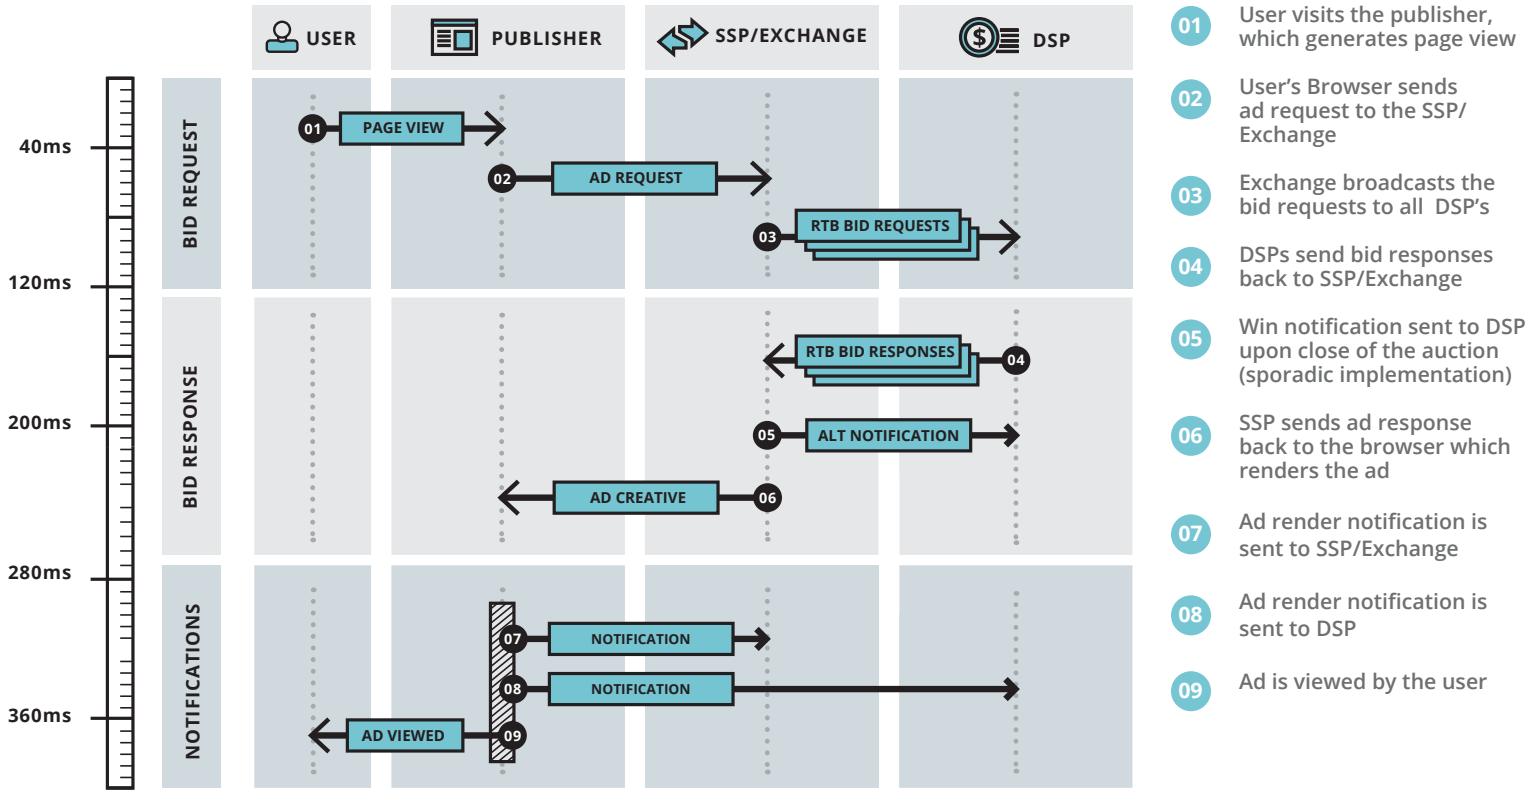

FIGURE 03 - NORMAL OPEN RTB FLOW

COPYRIGHT © 2015 PIXALATE

WHAT IS OPEN RTB'S IMPLEMENTATION ISSUE?

The OpenRTB v2.3 specification does not define a maximum period of time between generation of a bid response and the subsequent impression notification. Neither does it define a mechanism within the bid request for the DSP to supply a bid-specific time limit. If the DSP and/or ad exchange do not go beyond the specification on their own and implement time limits themselves, they are open to delayed bursts of fraudulent impressions.

In other words, even if the ad markup is rendered after two hours of bid response, it may still be counted as a valid impression by the bidding engine.

This vulnerability allows fraudsters to conceal the true status of an ad transaction, which in turn causes the bidding engines to bid on more impressions per compromised host than originally intended. Figure 2 shows how the fraudster exploits this vulnerability:

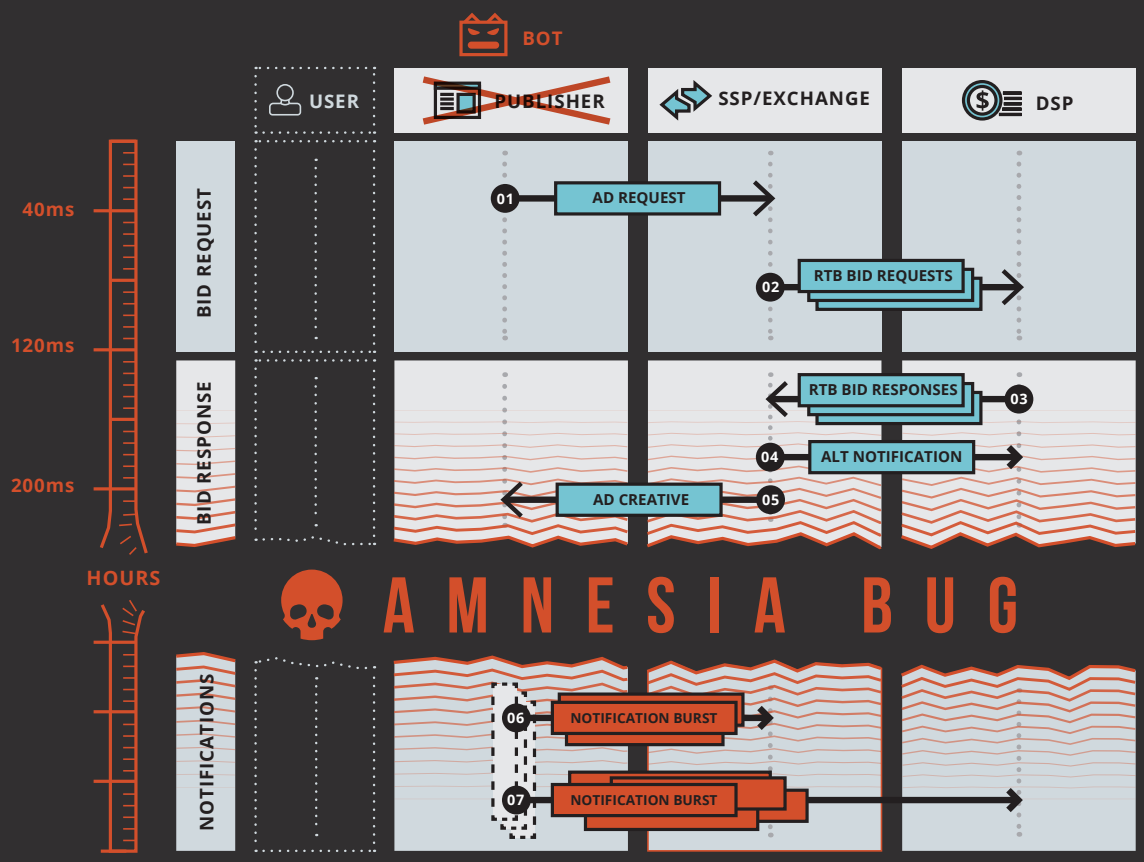

01 Bot instantiates a session on a browser and browser sends ad request to the SSP/Exchange

02 Exchange broadcasts the bid requests to a DSP

03 DSP send bid responses back to SSP/Exchange

04 Exchange optionally sends Win notification to DSP upon close of the auction

05 SSP/Exchange sends ad response back to the browser

06 Bot does not render the ad markup (instead hoards it) for a time period ranging from a few seconds to a few hours and repeats steps 1-7 multiple times

07 Bot replays the ad markups (it was hoarding) instantaneously and Ad render notifications get sent to SSP/Exchange

08 Ad render notifications get sent to DSP

FIGURE 04 - BOTNET IMPACTED OPEN RTB FLOW

COPYRIGHT © 2015 PIXALATE

HOW IS THE VULNERABILITY BEING EXPLOITED?

Xindi is installed via popular methods such as drive-by downloads, malware, and phishing attacks using social-engineering tactics, including but not limited to impersonating organizations.

Once Xindi corrupts a machine, it initiates an ad request to the ad server by visiting the domain.com, which loads a page with ad tags shown below:

```
<iframe width="300px"
height="250px" frameborder="0"
scrolling="no"
src="https://adserver.net/
getAd;myd=ad;sz=300x250;tile=1;dc_
ref=https%3A%2F%2FdomainToMask.
com%2Ffinance;domain=www.
domainToMask.
com;ord=892484284?"></iframe>
```

This page generates fraudulent ad requests to the ad exchange, and the ad exchange broadcasts bid requests to partner DSPs. Once the ad exchange receives the bid responses, the auction is completed. The ad exchange adds the DSP notification URL as well as its own beacon to the ad markup and passes it back to the browser as the ad response.

Once the ad response is received by Xindi, it will not render the ad markup but instead will hoard it for a few seconds to a few hours. Xindi will then play all the ad markups in a burst, instantaneously. Only at this point in time will the Ad Exchange as well as the winning DSP be notified of winning the auction.

```
<img src="http://ad.dspServer.com/
s=862230/size=728x90/cpm=1.52/
auction_id=123454321/" type="text/
javascript"></img>
```


COPYRIGHT © 2015 PIXALATE

XINDI BOTNET IMPACT

COPYRIGHT © 2015 PIXALATE

WHAT ARE THE TRENDS AND IMPACT?

Patterns of traffic resembling this botnet first appeared on October 29, 2014. Since then, follow-up attacks were observed in December 2014, March 2015 and August 2015. Overall, the volume of traffic exploiting this vulnerability has been following a steep curve since August 2015. The last attack, which occurred in August 2015, led to up to billions of fake impressions. The majority of the impact (i.e., more than 90 percent) was on campaigns targeting U.S. audiences.

Since its first appearance in October 2014, Xindi's footprint has continued to increase; it currently comprises more than 6 million machines, the majority of which are in the networks of enterprises and educational institutes. Based on its study, Pixalate currently estimates Xindi botnet will increase fraud by 50 percent and cost marketers more than $3 billion by the end of 2016.

> "THE LAST ATTACK, WHICH OCCURRED IN AUGUST 2015, LED TO UP TO BILLIONS OF FAKE IMPRESSIONS"

NUMBER OF UNIQUE MACHINES

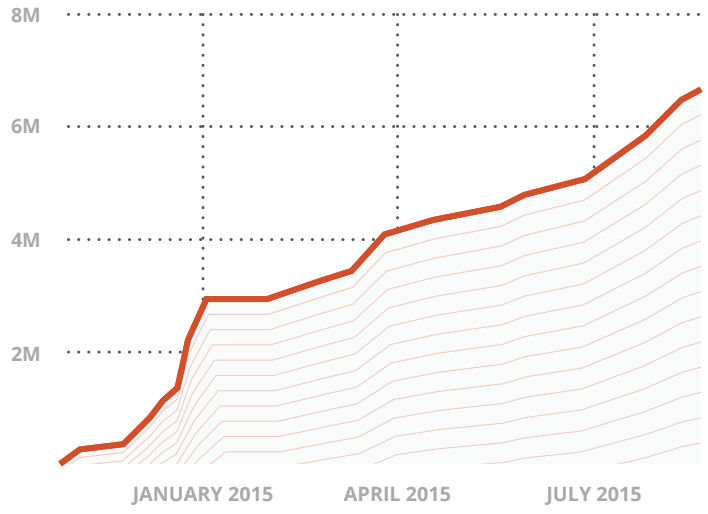

FIGURE 05 - NUMBER OF UNIQUE MACHINES

FRAUDULENT IMPRESSIONS

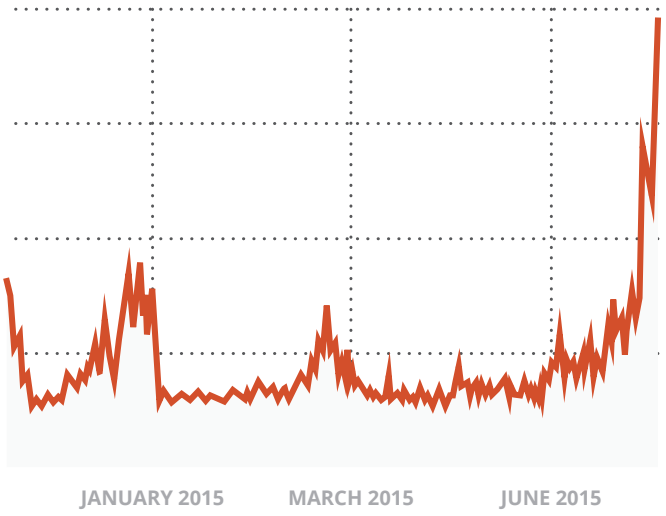

FIGURE 06 - FRAUDULENT IMPRESSIONS

COPYRIGHT © 2015 PIXALATE

WHAT IS THE GEOGRAPHIC DISTRIBUTION OF XINDIBOT?

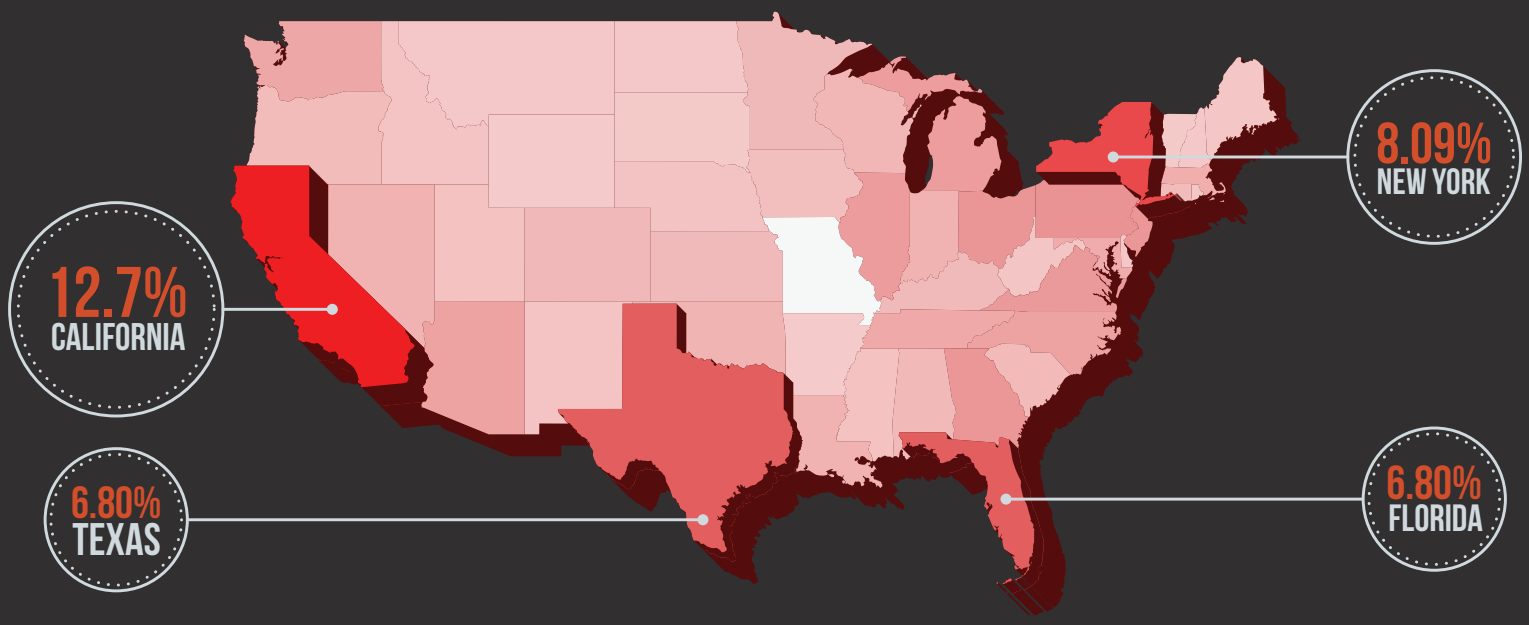

FIGURE 07 - XINDI BOTNET GEO DISTRIBUTION

WHAT IS THE OS DISTRIBUTION OF XINDIBOT?

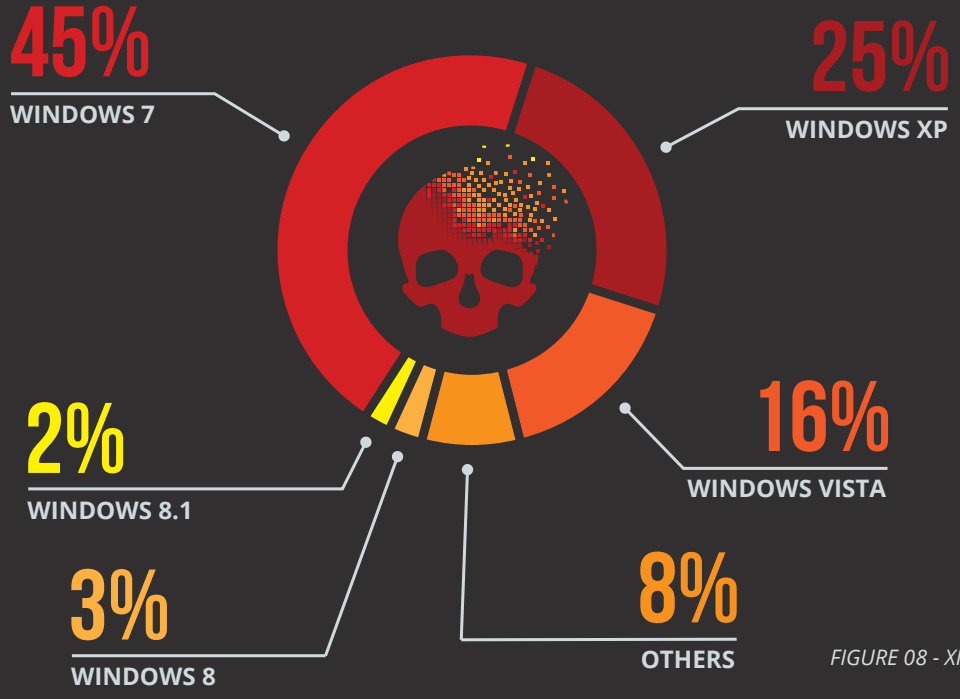

FIGURE 08 - XINDI BOTNET OS DISTRIBUTION

COPYRIGHT © 2015 PIXALATE

WHAT ORGANIZATIONS ARE AT RISK?

As part of this study, Pixalate has generated security risk scores for over 80,000 companies and universities by monitoring honeypot and malware sites that are being visited by botnet-infected endpoints.

Leveraging an advertising channel as a sinkhole, Pixalate analyzed a company's Internet traffic from the outside and built its profile by correlating the traffic data to its IP addresses belonging to these organizations.

For the purposes of this study, Pixalate has created the following security index representing a company's overall risk score.

METHODOLOGY

The following Security Index is a weighted analysis of three factors: malware Risk Score, egress traffic risk score and endpoint security risk score.

- **Malware Risk Score:**
  The number of visits to malware infected sites and command and control centers. The greater the number of visits, the higher the company's risk score.
- **Egress Traffic Risk Score:**
  This measures the hygiene of the enterprise's outgoing traffic.
- **Endpoint Security Risk Score:**
  Outlier behavior detection using machine metadata such as active or inactive plugins installed on an organization's endpoint.

EDUCATION

RANK	ORGANIZATION	OVERALL RISK SCORE	OVERALL RISK LEVEL	MALWARE RISK SCORE	MALWARE RISK LEVEL	EGRESS TRAFFIC RISK	EGRESS TRAFFIC RISK LEVEL	ENDPOINT SECURITY RISK SCORE	ENDPOINT SECURITY RISK
1	California State University - Office of the Chancellor	97.5	Critical	96.2	Critical	87.3	Critical	72.6	Moderate
2	Columbia University	83.6	High	84.7	High	76.1	Moderate	70.7	Moderate
3	University of Maryland	82.1	High	82.9	High	77.6	Moderate	71.0	Moderate
4	University of California - Office of the President	82.1	High	83.1	High	75.0	Moderate	70.6	Moderate
5	Kennesaw State University	80.3	High	80.3	High	89.3	High	71.4	Moderate
6	Texas A&M University	79.5	Moderate	80.1	High	76.7	Moderate	70.5	Moderate
7	Rutgers University	79.0	Moderate	79.6	Moderate	76.1	Moderate	70.8	Moderate
8	New York University	78.5	Moderate	79.1	Moderate	74.5	Moderate	70.4	Moderate
9	University of California, Los Angeles	78.3	Moderate	78.9	Moderate	75.5	Moderate	70.6	Moderate
10	Boston University	68.1	Moderate	78.5	Moderate	76.2	Moderate	72.5	Moderate

COPYRIGHT © 2015 PIXALATE

WHAT ORGANIZATIONS ARE AT RISK? (CONT.)

FINANCE

RANK	ORGANIZATION	OVERALL RISK SCORE	OVERALL RISK LEVEL	MALWARE RISK SCORE	MALWARE RISK LEVEL	EGRESS TRAFFIC RISK	EGRESS TRAFFIC RISK LEVEL	ENDPOINT SECURITY RISK SCORE	ENDPOINT SECURITY RISK
1	Wells Fargo & Company	96.1	Critical	99.4	Critical	92.0	Critical	94.3	Critical
2	Citigroup	95.8	Critical	96.1	Critical	86.5	Critical	89.0	Critical
3	LPL Financial	87.5	Critical	90.5	Critical	61.6	Moderate	60.5	Moderate
4	Bloomberg Financial Market	81.5	High	83.2	High	67.6	Moderate	64.2	Moderate
5	Morgan Stanley Group	81.4	High	81.1	High	89.4	Critical	78.7	High
6	Deutsche Bank	80.3	High	82.0	High	65.3	Moderate	64.0	Moderate
7	Harris Trust & Savings Bank	78.9	High	80.7	High	64.4	Moderate	61.3	Moderate
8	Mellon Bank	73.9	High	75.1	High	65.6	Moderate	60.8	Moderate
9	Compass Bank	73.8	High	75.2	High	62.7	Moderate	61.2	Moderate
10	Bank of America	72.6	High	73.5	High	67.2	Moderate	63.3	Moderate

FORTUNE 500

RANK	ORGANIZATION	OVERALL RISK SCORE	OVERALL RISK LEVEL	MALWARE RISK SCORE	MALWARE RISK LEVEL	EGRESS TRAFFIC RISK	EGRESS TRAFFIC RISK LEVEL	ENDPOINT SECURITY RISK SCORE	ENDPOINT SECURITY RISK
1	General Motors	99.5	Critical	99.7	Critical	91.0	Critical	89.0	Critical
2	Lowe's Companies	99.1	Critical	99.5	Critical	90.4	Critical	92.4	Critical
3	Marriott International	98.2	Critical	99.4	Critical	85.3	Critical	85.3	Critical
4	Wells Fargo & Company	96.1	Critical	99.4	Critical	92.0	Critical	94.3	Critical
5	Citigroup	95.8	Critical	96.1	Critical	86.5	Critical	89.0	Critical
6	CenturyLink	91.2	Critical	90.7	Critical	90.3	Critical	99.8	Critical
7	Charter Communications	89.3	Critical	89.7	Critical	91.2	Critical	78.8	High
8	Morgan Stanley Group	81.4	High	81.1	High	89.4	Critical	78.7	High
9	Bank of America	72.6	High	73.5	High	67.2	Moderate	63.3	Moderate
10	Fidelity National Financial	70.5	High	71.2	High	67.1	Moderate	61.5	Moderate

COPYRIGHT © 2015 PIXALATE

WHAT IS THE PROJECTED LOSS?

In advertising campaigns where Xindi botnet activity was observed, fraud went up by 300 percent. These campaigns tend to have high viewability, in the range of 85-90 percent. Given that the impressions are released at same time, quality gates such as frequency cap are not able to block Xindi and prevent fraud. The study shows that by the end of 2016 the advertising industry will lose up to $3 billion due to the Xindi botnet.

> "BY 2016 THE ADVERTISING INDUSTRY WILL LOSE UP TO $3 BILLION DOLLARS DUE TO XINDI BOTNET."

BOTNET REVENUE IMPACT

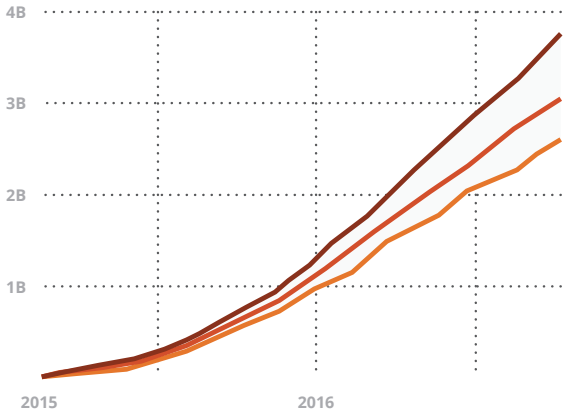

	IMPACT	CPM
MAX	3.6 BILLION	$1.50
AVG.	2.9 BILLION	$1.20
MIN	2.4 BILLION	$1.00

FIGURE 09 - XINDI BOTNET PROJECTED REVENUE LOSS

If your company is involved in buying and/or selling digital advertising programmatically, you have most likely been affected by the Amnesia Bug.

Most notably, if your company is a supply-side platform (SSP), a demand-side platform (DSP), an advertising network, an advertising exchange or a trading desk using the OpenRTB protocol, it has most likely been affected. Agencies and brands employing an affected vendor have been impacted indirectly.

COPYRIGHT © 2015 PIXALATE

HOW IS THE VULNERABILITY BEING EXPLOITED?

When enough requests have been generated as described above, the Xindi botnet will continue the same process with a different domain, after changing its UA string in order to appear as a different user. The browser cookies are also deleted after being used for a domain, although in some cases cookies are not deleted while generating the request for other domains. In general, the deletion of browser cookies after each domain visit makes it harder to uncover any underlying fraudulent activities by an IP.

FIGURE 10 - COOKIE AND USER AGENT SPOOFING

COPYRIGHT © 2015 PIXALATE

WHICH ADVERTISERS ARE AT RISK?

Pixalate observed many cases of affected campaigns in the wild. for advertisers such as **Home Depot, Uber, Honda, Pandora and Monster**. According to the data, the following are the top advertisers that are being exploited using this vulnerability.

TOP AFFECTED ADVERTISERS:

- Home Depot
- Uber
- Marie Callender's
- McDonald's
- Honda
- Pandora
- Cricket
- Gerber Live
- Monster
- Verizon
- Nissan

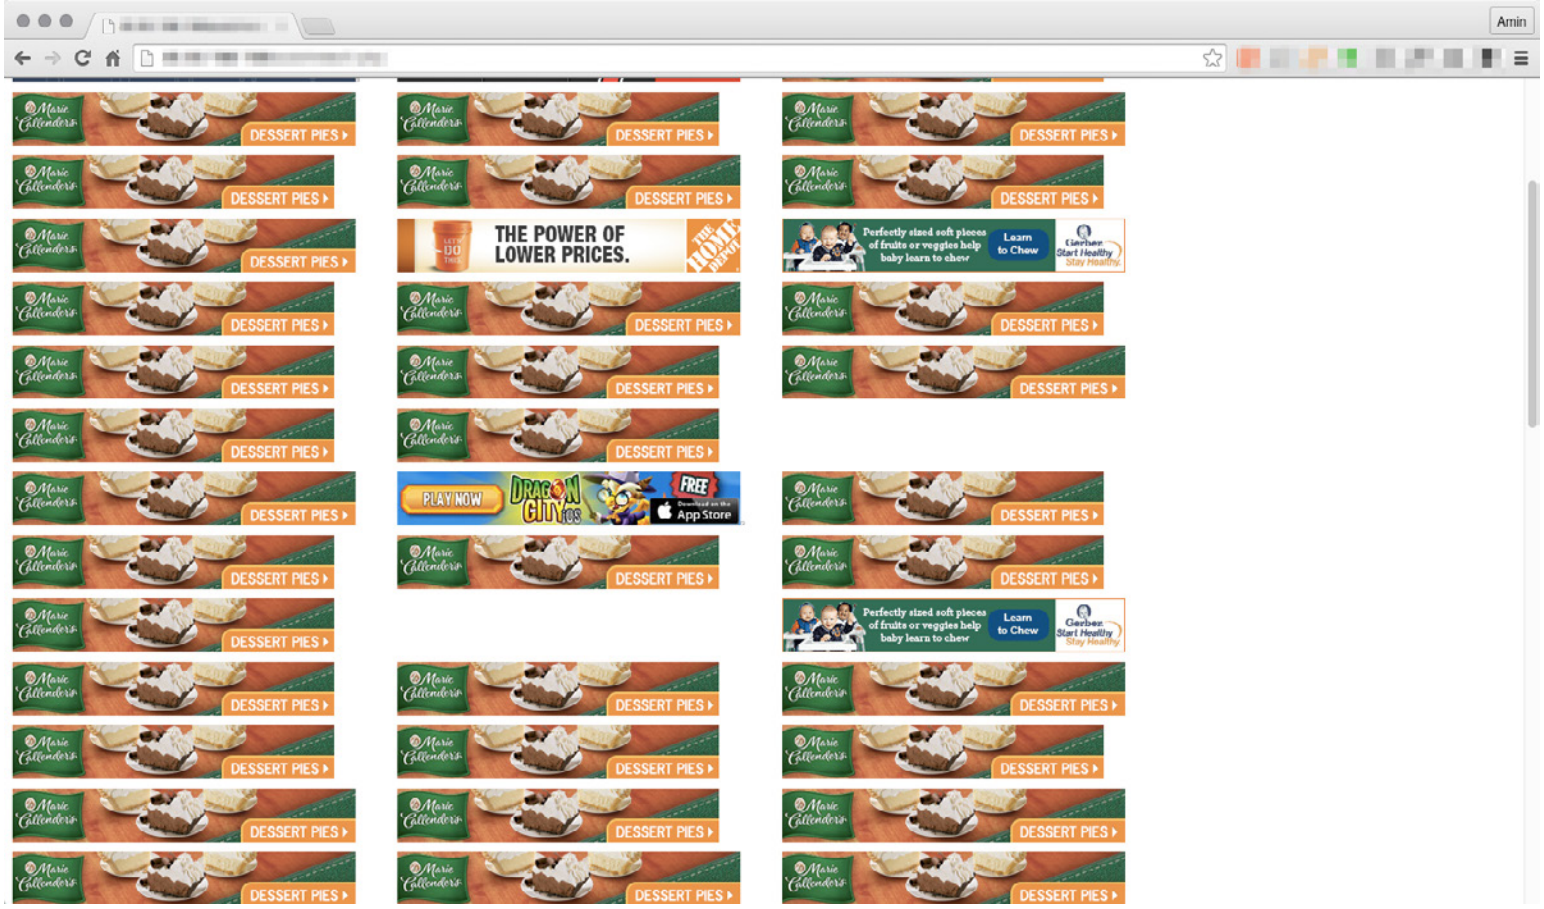

FIGURE 11 - AD FRAUD SCREENSHOT

COPYRIGHT © 2015 PIXALATE

WHICH ADVERTISERS ARE AFFECTED? (CONT.)

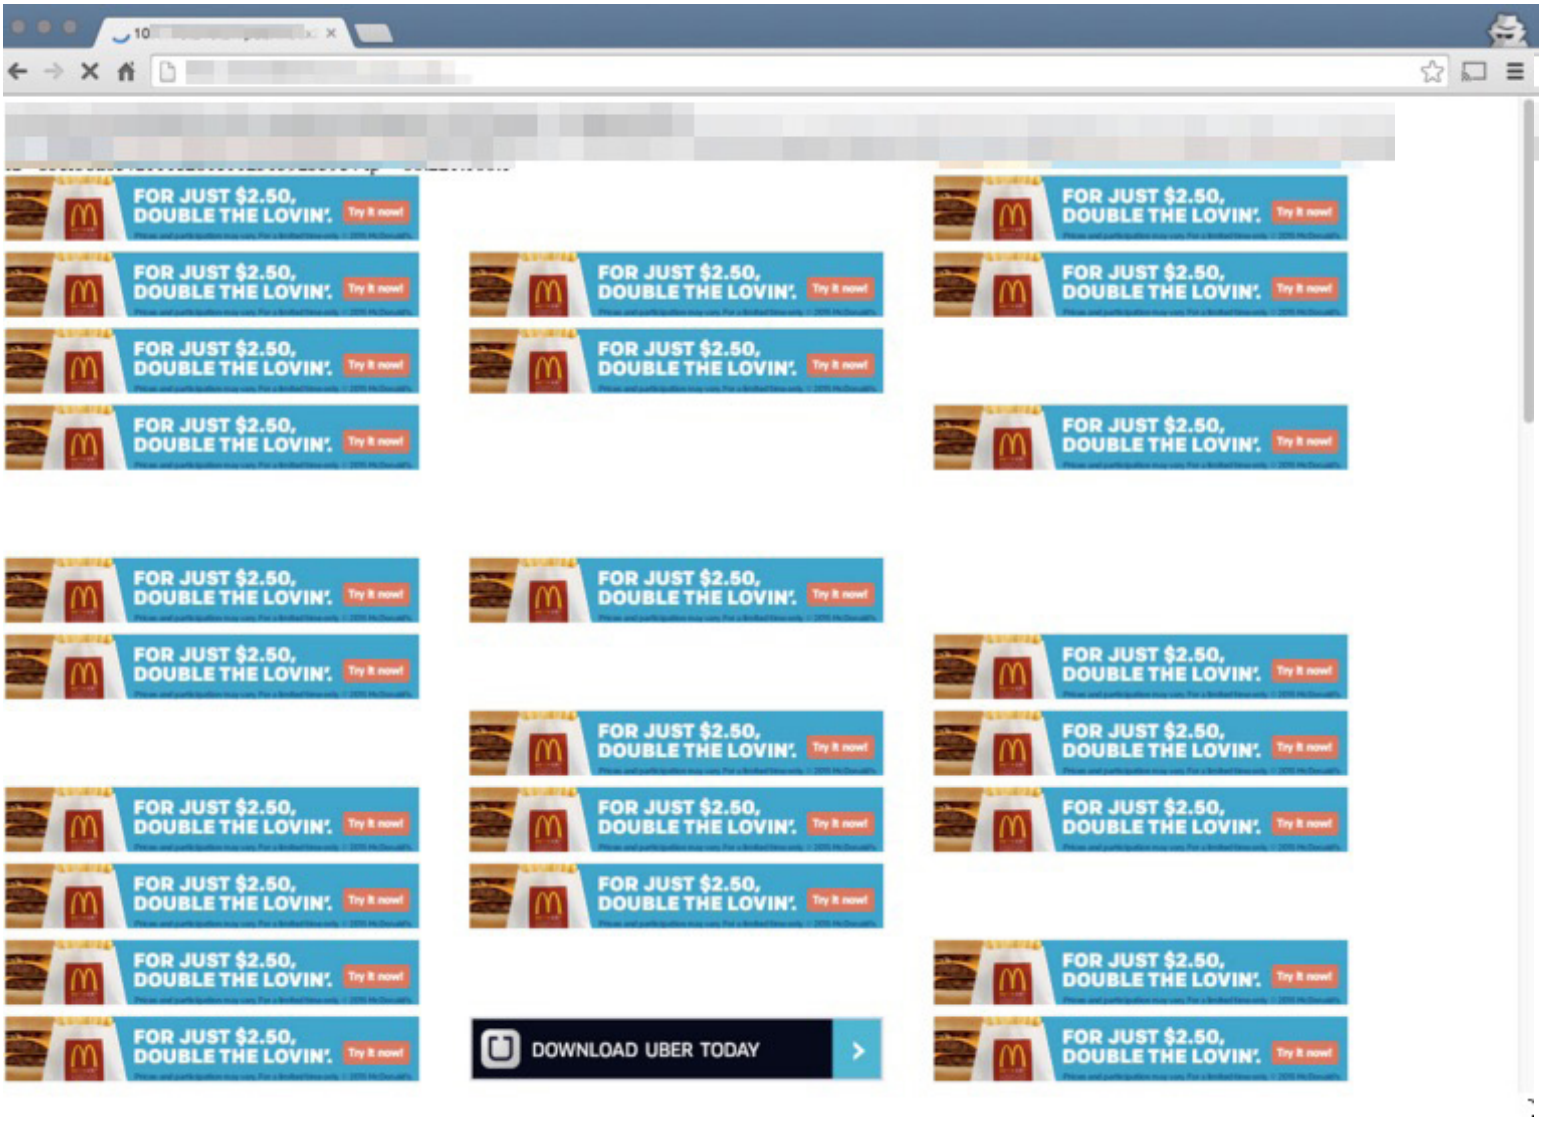

FIGURE 12 - AD FRAUD SCREENSHOT

COPYRIGHT © 2015 PIXALATE

FIX & MITIGATION

COPYRIGHT © 2015 PIXALATE

HOW CAN ADVERTISERS PROTECT THEMSELVES?

Advertising vendors can mitigate this issue by ensuring a reasonable bid timeout is in place in the auction algorithm. This will ensure that impressions generated after a certain time period not be accepted as valid - hence non-billable.

Long-term mitigation requires the Interactive Advertising Bureau (IAB) to update the protocol to account for reasonable bid timeouts, emphasize mandatory notifications to the winning party and encourage Exchanges/SSPs and DSPs to implement these remedies in order to cut fraud out of the ecosystem and improve quality.

Pixalate's security team is recommending a fix in order to make the internet a safer place for everyone. However, we can't guarantee this recommendation will work or is appropriate for your environment. You should assess your own vendors and consult experts as necessary to address your specific situation. Your use of any information in this document is at your own risk. All information is provided "AS IS" without any warranty of any kind. We expressly disclaim any liability for any loss or damages resulting directly or indirectly from your use of the information.

HOW CAN I KNOW IF MY ORGANIZATION IS COMPROMISED?

Pixalate is releasing the complete list of infected IP addresses, which have been exploiting this vulnerability. Enterprise and Universities Network Security teams can request access to the IP addresses belonging to their organizations. For further details, please contact Pixalate at **cisa@pixalate.com**

HOW CAN BRANDS PROTECT THEIR DOLLARS FROM AMNESIA?

If a vendor manages your media buys, contact the vendor's support team and check how its staff has employed safeguards (including making public disclosures) against the Amnesia Bug. You can also check Pixalate's Affected Vendors list for more updates.

COPYRIGHT © 2015 PIXALATE

CONNECT WITH US

join the conversation on the web:

#AMNESIABUG

GET IN TOUCH

info@pixalate.com

FIND US ONLINE

pixalate is an Enterprise Security and Analytics Platform that provides real-time fraud protection and threat intelligence to Marketers and Fortune 500 companies. Pixalate has disrupted the traditional security model by developing technology that infiltrates the attackers' network revealing the "mothership" of botnets and infected machines worldwide.

For more information visit
PIXALATE.COM

Vulnerability Reference - CVE-2015-7266
Open RTB 2.3 Specification - OpenRTB 2.3 specification - IAB.net

DISCLAIMER. Pixalate's analysis of how organizations are exposed to risk from advertising fraud, as well as its identification of specific organizations at various levels of risk, reflect Pixalate's opinions with respect to numerous factors. The scores or ratings in each category reflect Pixalate's opinions regarding the potential exposure of a particular organization, but is not intended to suggest that a particular organization's computer system has in fact been infected. That is a matter within the domain of each organization's cyber security function. The ratings and rankings reflect Pixalate's opinions regarding the relative risk identified for each firm with respect to the specified factors. Pixalate's opinions are just that – opinions, not facts or guarantees – but they are grounded in its proprietary technology and analytics, which Pixalate is continuously evaluating and updating.

COPYRIGHT © 2015 PIXALATE