sandvine®

Intelligent Broadband Networks

2017 Global Internet Phenomena

SPOTLIGHT: ZERO-RATING FRAUD

sandvine®
Intelligent Broadband Networks

Global Internet Phenomena Spotlight: Zero-Rating Fraud

Introduction

Communications service providers (CSPs) are under tremendous competitive pressures due to flat or declining average revenue per user (ARPU) and slowing or non-existent subscriber growth. Accordingly, CSPs must look to grow by attracting new subscribers from competitors' networks. With this fiercely competitive environment, it is imperative for CSPs to launch new and differentiated services that both build loyalty with existing customers and entice potential subscribers to switch providers.

One proven, effective way that CSPs differentiate services and stand out from the competition is by zero-rating applications. Zero-rating is a data offering that enables unlimited usage of one or many applications, services, or websites. This approach differs from data plans where a customer pays a fixed amount (prepaid or postpaid) for a specific access speed and volume quota (e.g., a particular number of megabytes or gigabytes, over a time period). Zero-rating also includes both paid offerings (e.g., a fixed price for Unlimited Social Networking or Music Streaming), free offerings (e.g., Internet as a Public Service, such as Free Basics by Facebook and others), and the CSP's web portal and self-care applications.

Enticed by the potential of receiving unlimited data for minimal (or no) cost, subscribers have a strong incentive to engage in circumventing behavior. This behavior is achieved by disguising data traffic to look like zero-rated content in order to circumvent a CSP's network charging rules. A variety of applications to assist with such circumvention exist online, some of them associated or derived from services created to enable users bypass broadly deployed Internet filtering schemes[1]. For instance, a subscriber could purchase a zero-rated Facebook plan and use one of the either free or commercially available techniques that make all data traffic look like it was associated with Facebook. Should these techniques become widespread, then the result could significantly prevent CSPs ability to suit the needs of individual subscribers with more tailored plans.

This report examines a real-world example of how one tier-1 mobile operator experienced significant revenue leakage on their 2G and 3G networks by subscribers using HTTP Header Injection based applications to exploit zero-rating of traffic at two domains: the CSP's website and its self-care portal.

It is important to state that not all CSPs have outlined specific Terms and Conditions (T&Cs) to deal with the potential for Zero Rating abuse meaning that labeling these types of circumvention techniques as fraudulent could be controversial. In the context of this report the word fraud and its derivatives are used to denote any attempt to circumvent application based Zero Rating plans by using circumvention techniques that allow the use of applications not included in the Zero Rating list.

[1] Learn more here: https://en.wikipedia.org/wiki/Internet_censorship_circumvention#Proxy_websites

HTTP Header Injection

There are a number of techniques that can be used to take advantage of zero-rating services to commit billing fraud, and the specific process examined in this report will focus on HTTP Header Injection.

HTTP is one of the Internet's key communication languages and enables communications between clients (browsers) and servers (websites). The protocol uses a request-response communication method where a client submits an HTTP request to a server, and the server returns a response to the client. Header fields are included in each HTTP request and provide metadata about the request.

To enable zero-rating for particular websites, many Policy Control Enforcement Functions (PCEF) or Traffic Detection Functions (TDF) solutions look for and rely uniquely upon specific information within the HTTP header to determine whether or not the website should be zero-rated or not. For example, the PCEF or TDF may rely uniquely on the HTTP Host header field found in HTTP transactions.

Accordingly, if a fraudster examined traffic captures of zero-rated websites and compared them to non-zero-rated websites, they could discover the specific HTTP Host headers that enable zero-rating. Once the fraudulent user has this information, they could utilize one of many client applications to manipulate the HTTP Host headers and replace its content to trigger the PCEF/TDF to zero-rate the communication. This action tricks the system into giving the fraudulent user unlimited Internet browsing.

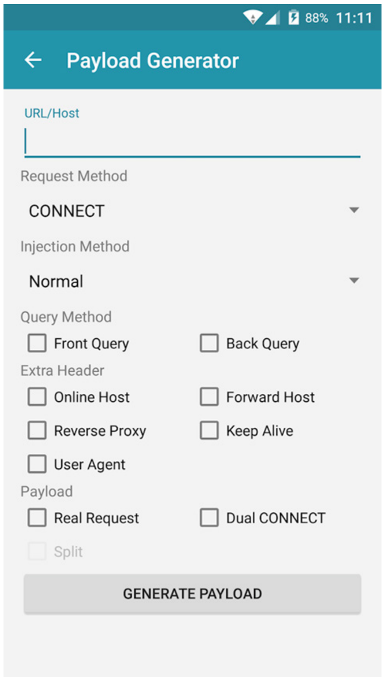

URL/Host

Request Method: CONNECT

Injection Method: Normal

Query Method

□ Front Query □ Back Query

Extra Header

□ Online Host □ Forward Host

□ Reverse Proxy □ Keep Alive

□ User Agent

Payload

□ Real Request □ Dual CONNECT

□ Split

Figure 1 - Screenshot from a leading HTTP Header Injection app

Figure 1 above is a screenshot of one popular app used to commit HTTP Header Injection fraud. While the description of the app in various app stores does not explicitly say this app and others like it can be used to commit fraud, subscribers will often find instructions on how to configure these applications on YouTube or online forums.

Interestingly, in order to use these apps to get free data, subscribers will often require paid data access to allow the HTTP Header Injection app to make an initial connection to the host. For post-paid subscribers that may mean they begin to use these services only after they have nearly exhausted their month quota, and for pre-paid users (the vast majority examined in this study), it means they may purchase a data package with a small usage quota, but then use HTTP Header Injection applications to get unlimited free data.

Additionally, although it was not prevalent on the network examined for this report, some HTTP Header Injection apps can even be configured to access Fraud as a Service (FaaS). In these cases, the subscriber pays a monthly subscription fee (lower than the cost of a data plan), to an illegitimate cloud service which will automatically manage the HTTP Header Injection for the subscriber. This service helps to ensure that subscribers don't have to worry about losing their "free" access as a CSP changes their zero-rating sites or attempt to mitigate fraudulent usage.

What Defines Fraudulent Usage?

Every operator will define fraudulent usage in a different way. For some operators, any attempt to misuse network resources may be classified as fraud, while others may allow a small amount of misuse before classifying it as fraud. The longstanding consensus in telecommunications is that some level of revenue leakage is inevitable to launch economically viable services. As long as CSPs can measure and contain such situations, Revenue Assurance teams can approve and manage its associated impacts.

For the operator in this study, through the use of network analytics applications, it was determined that any subscriber who uses more than 50 MB of data at the two most popular zero-rated domains in a month was suspected of engaging in fraudulent usage, and flagged for follow-up. This figure was determined by examining the amount of traffic that a visit to one of these domains typically generates, and how often subscribers typically visit them. The same network analytics lead to the conclusion that at this operator, a typical visit to the zero-rated sites by a subscriber drives 5MB of usage, and they will visit the site on average 2-3 times each month.

While there might be some users who legitimately use more than 50MB of data at those domains, this threshold enabled the CSP to flag specific subscribers for further investigation as part of the standard Fraud Management and Revenue Assurance (FMRA) practices and procedures.

Findings

For this report, a segment of a tier-1 operator's 2G and 3G network was examined over a thirty-day period. During this period, 175,000 subscribers, or approximately 0.9% of total subscribers (as shown in Figure 1) were observed to have exceeded 50MB per month of traffic at the zero-rated operator services we studied: the CSP's website and its self-care portal.

The CSP examined also zero-rates multiple other URLs on the network (often more than ten at a time), but the data in this report focuses on just the traffic to the website and subscriber portal because they were the most popular, and because other URLs zero-rated by the operator are frequently changing.

For the thirty days examined, the suspected fraudulent subscribers consumed 140 TB of traffic through the two zero-rated services. For context, the 140TB represented 3.6% of the total usage from the subscribers on the 2G and 3G network.

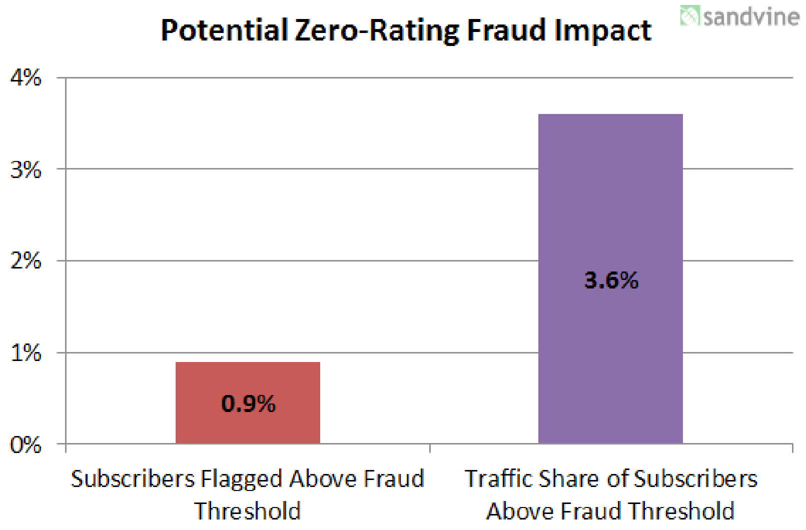

Figure 2 - Potential Zero-Rating Fraud Impact

Figure 3 below shows a comparison of mean monthly usage for both the network and those subscribers identified above the fraud threshold. On this network, mean monthly usage is 192MB, a figure lower than some may expect, but 3G devices use less data. Also, the figure is calculated based on the number of active devices on the network in a month, which likely includes subscribers who primarily use the LTE network (not the 2G or 3G network) but have temporarily dropped down to 3G at some point due to connection or coverage issues. Quite naturally, they would have low average usage on the 2G or 3G network.

Subscribers identified as being above the fraud threshold generated a monthly mean usage that was over 300% higher (805MB) than the mean monthly average for all subscribers on the 2G and 3G network.

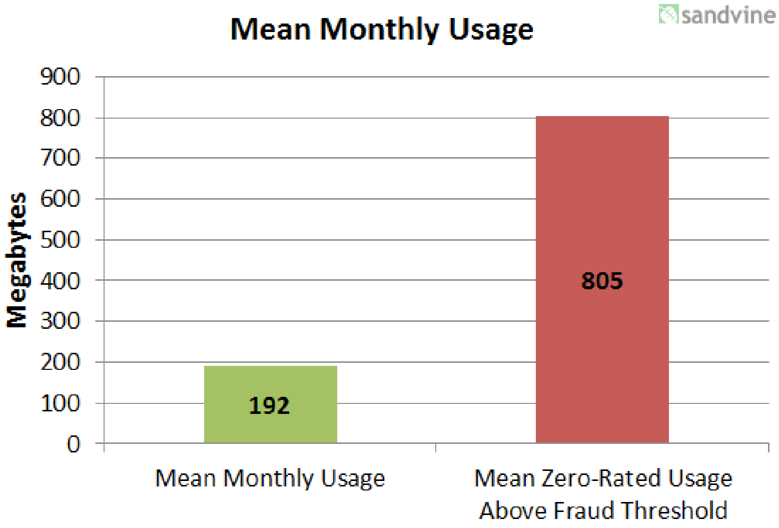

Figure 3 - Mean Monthly Usage

The final step in analyzing the potential fraud requires more advanced traffic categorization techniques that do not simply rely on HTTP Host header information. More advanced and robust categorization techniques allow for a breakdown, such as in Figure 4 below, of the 140TB of traffic observed for the identified subscribers. The usage data is very similar to typical usage for the region, which Sandvine publishes in our annual Global Internet Phenomena Reports. Popular social applications like Facebook, WhatsApp, and Instagram account for over one-third of the usage.

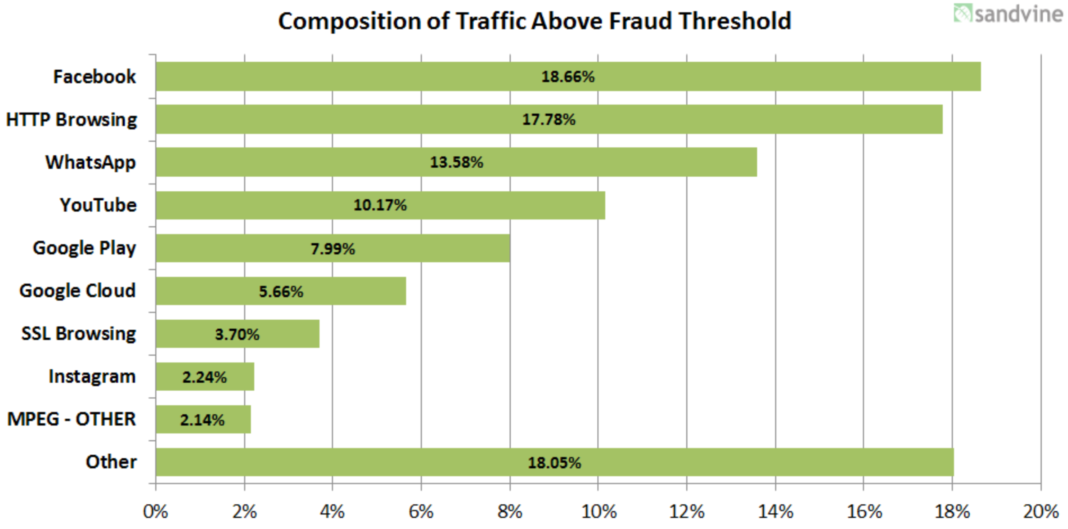

Figure 4 - Composition of Traffic above Fraud Threshold

This thorough analysis helps the CSP to accurately determine the nature of the suspicious activity and take the appropriate measures.

In addition to the 2G and 3G networks examined in this report, the CSP also operates an LTE network. With LTE offering speeds up to 25 times faster than 3G, the amount of data that can be defrauded via HTTP Header Injection is significantly higher. If we also assume that fraudulent users are more technically advanced than the typical user, it may be reasonable to assume that such users are more likely drawn to the LTE network. In combination, these factors could result in significantly more revenue loss on LTE networls than on the 2G and 3G networks studied in this report.

To support this theory, when the data in this report was initially shared with the CSP, they extrapolated that up to 10% of the usage on network traffic could be fraudulent and that the bulk of that fraud was occurring in the LTE network.

Calculating Potential Revenues Loss

Every operator calculates their capacity costs in different ways, and since that information is a tightly guarded secret within an operator's business, it is difficult for third parties to calculate how much fraudulent usage costs the operator.

What is possible, however, is to calculate the lost revenues from the data usage that was fraudulently zero-rated. To do this, Sandvine examined the five featured data plans available on the sample network's website and calculated that the average price for a gigabyte of usage per month is approximately $8 United States dollars (USD).

Extrapolating that figure out to the usage observed, it means that this operator could potentially be losing over $1.1 million USD a month ($8/GB x 140,000GB observed) from data usage in this portion of their 2G and 3G network. Taking that estimate even further, if the usage rates of this sample section of the network are equivalent to the entire network, the operator could be losing over $7 million USD a month from this single activity, and potentially far more in the LTE network due to the reasons mentioned earlier in this report.

The above calculations work under the assumption that all potentially fraudulent data usage would have been paid for by the subscriber, and that is obviously not the case. It is likely more data was used because it came at no cost to the subscriber. The calculations above however help to illustrate the potential magnitude of the issue and how CSPs could significantly boost revenue if they were able to convert fraudulent usage into paid usage.

Preventing Zero-Rating Fraud via HTTP Header Injection

The foundation of fraud management is having advanced traffic classification solutions in place so that most types of fraud can be averted and other signs spotted. To achieve this goal, a CSP must work with a best-of-breed PCEF/TDF solution that has experience in not only dealing with the issues of today, but that understands the ever-changing landscape of fraud attempts. Although many PGW and GGSNs offer rudimentary traffic identification capabilities, their technologies will likely prove insufficient to detect and prevent zero-rated fraud.

In the example discussed in this report, the operator leveraged several capabilities of Sandvine's Network Policy Control platform to gain unprecedented visibility of suspicious activity in its network and subsequently manage, and prevent future zero-rating fraud attempts.

Additional Reading

If you are interested in learning more about preventing zero-rated fraud, Sandvine has published a white paper that compliments this report entitled “Considerations and Best Practices for Zero-Rated Fraud Prevention.”[2]

The table below provides a high-level summary of that paper’s conclusions on the four best practices CSPs should undertake to prevent fraud, with the paper itself delving into much more technical detail.

Best Practice	Description
Advanced Traffic Detection Techniques	• To identify fraudulent activity, a CSP requires a DPI or gateway vendor that utilizes an advanced traffic detection capabilities like application fingerprinting. Application fingerprinting discovers fraud by comparing application data traffic to historical norms
Practical Policy Enforcement Rules	• CSPs should rely on their PCEF or Gateway charging vendor for specific instructions on how to logically code zero-rated policy • Enforcement instructions should be provided to CSPs in traffic classification profiles • Enforcement options should be configurable depending on the likelihood of fraud • e.g., Flag and measure fraud, notify customer and block traffic
Reporting Metrics that Measure Total Fraud	• To adequately measure fraud, a CSP must flag and report on all fraudulent activity • Fraud reports should include the following information: • The application used to enable the fraud (if applicable) • The subscriber using the fraud technique • The zero-rated application and/or plan being defrauded • The amount of fraudulent data used
Customer Communication	• Zero-rated messaging • Zero-rating plans must be transparent and easily understood by subscribers • Customers need to understand what data is and isn’t included in their zero-rated offering • They also need to understand what activities are and are not permitted under the zero-rated plan (e.g., Fair Use Policy) • A real-time notification system to communicate with subscribers • Notifying the customer is critical when fraud is detected • The notification should alert the user that fraud has been detected on their zero-rated plan, explain the corrective action, and link them to the CSP’s Fair Use Policy

[2] The whitepaper can be downloaded here: https://www.sandvine.com/resources/whitepapers/considerations-and-best-practices-for-zero-rated-fraud-prevention.html

Headquarters
Sandvine Incorporated ULC
Waterloo, Ontario Canada
Phone: +1 519 880 2600
Email: sales@sandvine.com

sandvine®
Intelligent Broadband Networks

www.sandvine.com
To learn more about the latest Internet trends: www.internetphenomena.com

Revision: 2017-04-18