



Zero-Rated Fraud Prevention:

Considerations and Best Practices

CONTENTS

Executive Summary	1
Introduction to Zero-Rated Fraud	1
Zero-Rated Fraud Techniques	2
HTTP Header Injection	2
Domain Fronting	2
DNS Spoofing	4
Best Practices for Preventing Zero-Rated Fraud	4
Traffic Classification	4
Policy Enforcement	7
Reporting Metrics	8
Subscriber Engagement	9
Conclusions	10
Additional Resources	11

EXECUTIVE SUMMARY

As Internet subscriber growth slows, the competition amongst communications service providers (CSPs) for new customers has intensified. Because of the increased competition, CSPs are developing new and differentiated services that enable them to stand out from the competition. Innovative new Internet services are more important than ever for both building loyalty with existing customers and for enticing potential subscribers to switch providers.

One of the most popular techniques that CSPs are using to differentiate their offerings is zero-rating. Zero-rating enables unlimited usage of an application (or many applications) for a fixed price.

However, before a CSP implements a zero-rated offering, they must understand the risks and remedies associated with zero-rated fraud.

To prevent zero-rated fraud and minimize revenue leakage, Sandvine recommends the following best practices:

- 1 Rely only upon advanced traffic classification technologies
- 2 Apply practical policy enforcement
- 3 Leverage reporting metrics that flag and measure fraud
- 4 Implement a notification system to communicate with subscribers

By adhering to these principles, CSPs will be able to launch successful zero-rated offerings that prevent fraudulent activities.

INTRODUCTION TO ZERO-RATED FRAUD

Communications service providers (CSPs) are under a tremendous amount of pressure due to flat or declining average revenue per user (ARPU) and slowing or non-existent subscriber growth. This pressure creates intense competition amongst CSPs because there aren't as many net new data customers as there were in the past. 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 Internet services that both build loyalty with existing customers and entice potential subscribers to switch providers.

One proven, effective way that CSPs differentiate their services and stand out from the competition is by zero-rating content. 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, perhaps over a time period). Zero-rating also includes both paid (e.g., Unlimited Social Networking or Music Streaming) and free offerings (e.g., Internet as a Public Service/Free Basics by Facebook).



Enticed by the potential of receiving unlimited data for minimal (or no) cost, deceitful subscribers have a strong incentive to engage in fraudulent behavior. This behavior is achieved by disguising data traffic to circumvent a CSP's charging rules. For instance, a subscriber willing to commit fraud could take advantage of the free data offered through the Free Basics by Facebook initiative and make all data traffic look like it was associated with Facebook. Should this technique become widespread, then the result would be significant revenue leakage for the CSP.

The purpose of this paper is to educate CSPs on the various types of fraud occurring in the market and to explain how to minimize the risk of revenue leakage associated with each type. More specifically, we're going to provide an overview on:

- Specific zero-rated fraud techniques, including HTTP Header Injection, Domain Fronting, and DNS Spoofing
- Approaches that minimize the impacts of zero-rated fraud, including Traffic Detection, Traffic Enforcement, Reporting and Auditing, and Communicating with Customers.

ZERO-RATED FRAUD TECHNIQUES

There are a number of techniques that can be used to take advantage of zero-rating services to commit billing fraud.

HTTP Header Injection

HyperText Transfer Protocol (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 deep packet inspection (DPI) or packet gateway solutions look for and rely upon specific information within the HTTP header to determine whether or not the website should be zero-rated or not.

For example, a zero-rated website (e.g., a CSP's customer portal) may contain text like "FreeZone" in the HTTP header. In contrast, "FreeZone" wouldn't be included in the HTTP header of websites that are not being zero-rated.

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 header (e.g., "FreeZone") that enables zero-rating. Once the fraudulent user has this information, they could utilize an HTTP injector application¹ to inject the "FreeZone" header into every HTTP request. This injection tricks the zero-rating system into giving the fraudulent user unlimited Internet browsing.

Due to the growing popularity of Internet as a Public Service programs (e.g., zero-rating access to essential websites), CSPs must be diligent in preventing HTTP header fraud.

Domain Fronting

Domain fronting is a masquerading technique that is typically used to circumvent Internet censorship by making traffic look like it's associated with a web domain that isn't restricted.

Although domain fronting wasn't initially intended to enable zero-rated fraud, the traffic masquerading techniques can achieve just that result. For example, a fraudster could disguise all of their Internet traffic to look like Facebook, and thereby take advantage of a zero-rated Facebook plan. To achieve this deception, domain fronting relies on content delivery networks (CDNs) that host multiple domains (websites).

¹ Such applications are readily available in any app store



CDNs host much of the web’s content in servers distributed around the world, typically as close to the subscriber access network as possible. A single CDN may host thousands of different websites and services (e.g., videos, software updates, etc.), even if the CDN itself is operated by a single company like Akamai, Cloud Front, Azure (Microsoft), or CloudFlare.

Because of the nature of CDNs, CSPs cannot simply block them, as that would unintentionally block many major websites, applications, and services. For example, facebook.com typically serves content from IP addresses owned by Facebook. However, in some emerging markets, Facebook content is served from domains like e239.akamai.net.

To enable domain fronting, most fraudulent subscribers utilize an application like Psiphon to route their Internet traffic to a CDN server. When the fraudulent user’s traffic reaches the CDN, it’s re-routed through a domain fronting server (e.g., Psiphon’s server) to its end destination². This re-routing process effectively masks the user’s traffic and makes it appear like all Internet traffic is coming from a legitimate website or application hosted on the CDN. Since major Internet players like Facebook distribute their content from CDNs, and since Facebook is included in many zero-rated Free Basics offerings, accurately identifying zero-rated traffic can be a significant challenge.

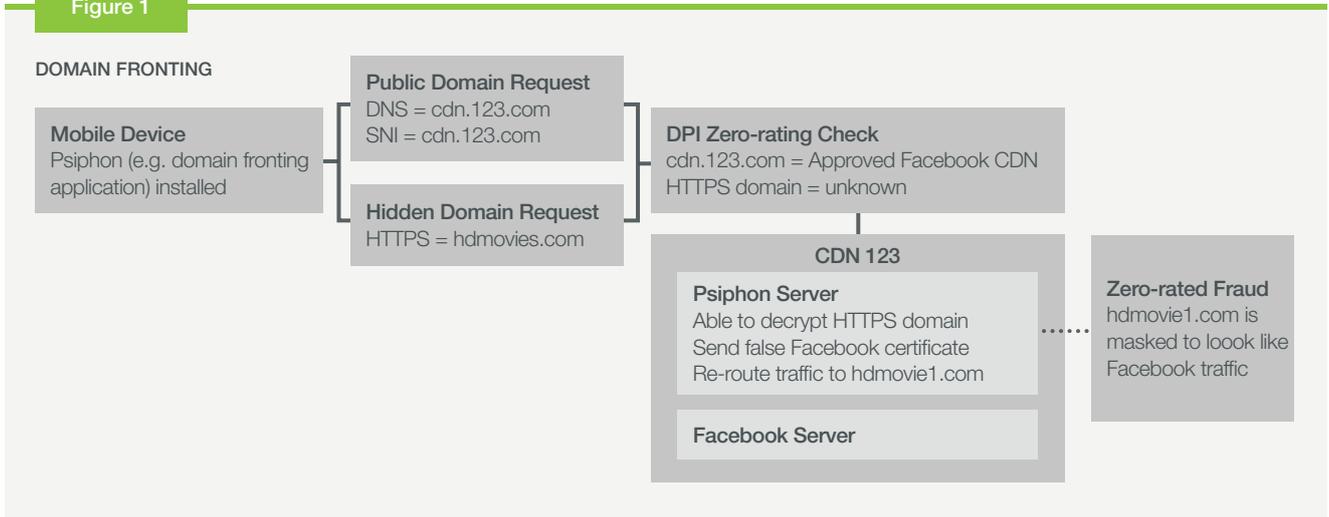
Through domain fronting, a fraudulent subscriber could potentially route all of their Internet traffic through Psiphon’s servers and make it look nearly identical to legitimate zero-rated application traffic – causing significant revenue leakage for a CSP.

Figure 1, below, shows the process for domain fronting. With normal Internet behavior, when a user visits a website, they send three domain name (e.g., website address) requests. These are sent over the DNS, TLS (more specifically the Server Name Indication or SNI field within TLS) and HTTP protocols. Generally, the three protocols include the same domain name to avoid confusion when searching for the location of the requested content. However, when a user is domain fronting, they purposefully include (typically done with an app like Psiphon) a different domain name in the HTTPS (the encrypted version of HTTP) protocol. Since the HTTPS domain name is encrypted, the zero-rating solution doesn’t see that the HTTPS domain name doesn’t match with the DNS and SNI fields, and only considers DNS and SNI when reviewing the traffic. This partial scan is unable to see the true nature of the Internet traffic and instead believes that the user is accessing legitimate (i.e., zero-rated) content hosted on a CDN. Once inside the CDN, the HTTPS domain name is decrypted by a domain fronting server (e.g., Psiphon’s server) and the traffic is re-routed to the hidden destination.

² Some users may set up their own server within a CDN to re-route traffic, but doing so takes significantly more effort than using a pre-made end-to-end service like Psiphon

This process enables fraudsters to trick most zero-rating solutions into determining that all Internet traffic is coming from a zero-rated service hosted in the CDN.

Figure 1





DNS Spoofing

The Domain Name System (DNS) is one of the many protocols that are a part of the Internet's infrastructure; in fact, DNS is often characterized as the Internet's phonebook because it provides the IP addresses that correspond to websites, Internet services, etc.

Since the DNS protocol is critical for delivering web content to users, blocking it to prevent fraud isn't an option. As a result, DNS traffic is trusted implicitly and rarely restricted³. A consequence of these characteristics is that DNS is an ideal protocol to circumvent inspection and masquerade data traffic.

DNS spoofing is another masquerading technique that disguises illegitimate traffic and makes it look like legitimate traffic. To enable DNS spoofing, a user utilizes an alternative DNS service configured to provide false DNS responses; these responses purposefully associate an incorrect IP address with a zero-rated website or service, and are observed and trusted by the zero-rating solution.

For example, if a zero-rating solution sees a fraudulent DNS response that lists mobile.twitter.com at IP address A.B.C.D, then the solution concludes that all traffic to and from A.B.C.D is Twitter.

If a fraudster has a zero-rated Twitter plan, then all he or she needs to do is use the fraudulent DNS service to trick the zero-rating solution into thinking that all traffic is Twitter.

It's important to note that not all alternative DNS services are used to enable zero-rated fraud. In fact, there are many legitimate reasons (e.g., faster speed, better reliability, enhanced security features, etc.) for a subscriber to use a public DNS service like Google Public DNS or OpenDNS. These legitimate third-party DNS services make the CSP's counter-fraud efforts much more challenging, because it's not a fair or correct conclusion that a third-party DNS is necessarily fraudulent.

BEST PRACTICES FOR PREVENTING ZERO-RATED FRAUD

To prevent zero-rated fraud and minimize revenue leakage, Sandvine recommends the following best practices:

- 1 Rely only upon advanced traffic classification technologies
- 2 Apply practical policy enforcement
- 3 Leverage reporting metrics that flag and measure fraud
- 4 Implement a notification system to communicate with subscribers

Traffic Classification

The foundation of fraud prevention is having advanced traffic classification technologies in place, so that the signs of fraud can be spotted.

To achieve this goal, a CSP must work with a best-of-breed traffic identification or DPI vendor. Although many GGSNs (or similar gateways) offer rudimentary traffic identification capabilities, their technologies – frankly – are insufficient to detect and prevent zero-rated fraud. This technical limitation makes sense, as DPI is not a core technology for a gateway, so identification techniques are more basic due to the limited processing power. Furthermore, since identifying encrypted applications and fraudulent techniques requires significant computing resources (e.g., processing power and memory), the accuracy difference between DPI and gateway vendors will become even larger as encryption and fraud techniques grow.

For simplicity, this paper identifies four high-level traffic detection techniques: deterministic, heuristic/machine learning, application signatures, and application fingerprints; an overview of each is included in Table 1.

³ Two notable exceptions include cyber security mechanisms that protect the DNS infrastructure, and DNS filtering technologies like those used in parental control and for regulatory compliance



However, we will spend the bulk of our efforts discussing Application Fingerprints, as they are the most advanced traffic identification method and the most relevant for fraud prevention. For an in-depth look at traffic identification and measurement, please review Sandvine's whitepaper, Identifying and Measuring Internet Traffic: Techniques and Considerations.⁴

Table 1

Method	Description
Deterministic	<ul style="list-style-type: none"> This traffic identification technique reviews port numbers, server names, IP addresses, URL addresses, byte patterns, cross-packet correlation, signatures etc. associated with the application to accurately identify it Combining multiple techniques creates an application signature; the more techniques combined, the stronger the signature Deterministic methods are the most accurate type of traffic identification because they achieve the lowest rates of false-positives and false-negatives Deterministic methods are also typically the fastest at identifying traffic (i.e., within the first few packets of a flow)
Heuristic/Machine Learning	<ul style="list-style-type: none"> Heuristics are used when a deterministic match or measurement is not available directly from the inspection of traffic By measuring the properties of traffic, conclusions can be reached that are sufficient to meet the immediate classification goals Machine Learning is a type of heuristic algorithm built on mathematical models By taking a set of known data and running it through a machine learning algorithm, correlations between the data set and the attributes associated with the traffic can be found Heuristics/Machine Learning techniques are often utilized as a second factor to provide more accurate traffic identification (e.g., YouTube SD vs. HD) and to identify fraud (e.g., data traffic doesn't match behavioral norms) In some instances, Heuristic/Machine Learning identification methods are not as quick to identify traffic as Deterministic methods because they must examine the traffic for a short period of time (e.g., 1-3 seconds) before accurately identifying it Heuristics/Machine Learning techniques are required to identify specific features within an application (e.g., voice call vs. instant message in WhatsApp) traffic should only be managed when congestion is present
Application Signature	<ul style="list-style-type: none"> A combination of multiple deterministic and heuristic/machine learning techniques are used to create an application signature Many applications are easy to detect and only require a small number of deterministic detection techniques in the application signature
Application Fingerprint	<ul style="list-style-type: none"> Application fingerprints build on application signatures by creating behavioral norms for applications Fraud is detected when behavioral norms deviate from the application fingerprint This is the most advanced type of traffic detection available and is critical for identifying fraud

Figure 2, below, presents the difference between an application signature and an application fingerprint. Clearly, an application fingerprint contains much more information than a typical application signature. By comparing observed behavior against this fingerprint, potential fraudulent activity is identified. So, even if fraudulent traffic appears legitimate upon inspection of IP address, host, port, and DNS, the fraud becomes apparent when the fingerprint is applied.

4 <https://www.sandvine.com/downloads/general/whitepapers/identifying-and-measuring-internet-traffic.pdf>

Figure 2

Application Signature vs. Application Fingerprint

Application Signature

- IP address = 1.23.456.78
- Host = examplesignature.com
- TCP Port = 1234
- DNS = 1.2.3.4.5

Application Fingerprint

- IP address = 1.23.456.78
- Host = examplesignature.com
- TCP Port = 1234
- DNS = 1.2.3.4.5
- Normal # of flows = 3-4
- Normal Bit Rate = 100MB/second
- Normal Up:Down Ratio = 1:10



To identify fraud, a solution must be able to collect and examine many more data points than are typically relied upon for traffic identification by simpler systems. These additional data points include behavioral norms for an application and are measured with heuristic/machine learning techniques, which are beyond the capabilities of most solutions. Examples of behavioral norms include (but aren't limited to):

- 1 The typical number of traffic flows (e.g., number of different servers populating content) that the application produces
- 2 The average traffic bitrate for the application
- 3 The median upstream to downstream (Up:Down) traffic ratio

If any of the above measurements significantly deviate from the historical norms built into an application fingerprint, then there is a high probability of fraud.

Consider a specific example: a fraudster is masquerading all of their traffic to look like Facebook traffic. How does application fingerprinting identify this fraud?

We know from detailed historical observation that Facebook typically utilizes 4-5 separate traffic flows (e.g., different servers) to power the user experience: a user's wall, online advertisements, Facebook videos, Messenger, and trending topics (see **Figure 3**, below).

However, if a subscriber was domain fronting and disguising additional traffic to look like Facebook, then only a single traffic flow would be recognized and it would have very different bitrate characteristics than usual Facebook traffic.

The difference between legitimate Facebook traffic and domain fronted traffic exists because domain fronting can only masquerade the traffic flow coming from a compromised CDN. As a result, even if fraudulent Facebook traffic is able to trick the deterministic measurements in an application signature (e.g., IP address and DNS match), the technique isn't able to match the behavioral norms in an application fingerprint (i.e., it doesn't have 4-5 traffic flows or the correct bitrate).

Figure 3

Normal Traffic Flow for Facebook

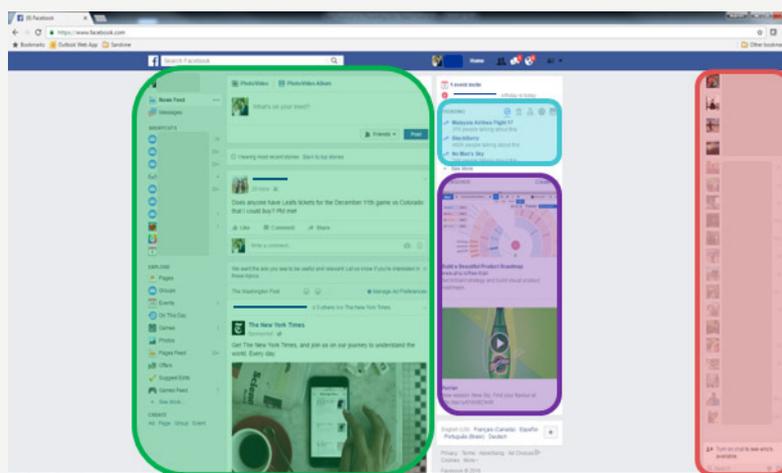
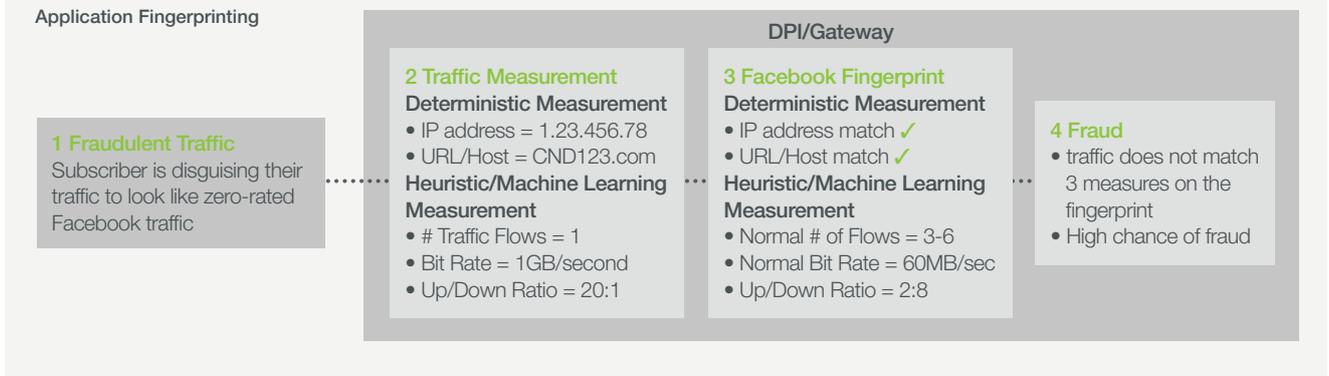


Figure 4 on the following page, highlights the application fingerprinting process. In this example, a user is trying to defraud a zero-rated plan and has successfully deceived deterministic measurements typically found in application signatures; however, through application fingerprinting, the traffic is clearly identified as fraudulent.



Figure 4



Since new applications are constantly being developed and existing applications are constantly adding new features, CSPs must select a vendor that issues timely updates to their traffic identification protocols (e.g., application fingerprints) to ensure that zero-rating and fraud detection remain accurate.

Policy Enforcement

The next key component to minimize the impact of zero-rated fraud is the ability to apply rational traffic enforcement policies. To correctly enforce policy, it is strongly recommended that CSPs work with vendors that have robust and up-to-date traffic classification profiles. Traffic classification profiles provide CSPs with specific instructions on how to enforce zero-rated traffic to minimize fraud. Much like application fingerprinting, traffic classification policies also require regular updates to ensure that enforcement policies remain accurate.

Please note that this paper specifically focuses on enforcement policies related to masquerading fraud. If you're interested in learning about how to enforce zero-rated application plans, please review Application Zero-Rating: Considerations and Best Practices for Revenue Assurance⁵.

Figure 5, below, highlights a simple zero-rating policy that is easily tricked by HTTP header fraud or DNS spoofing. A fraudulent user simply needs to inject the term "example.com" into the header or re-route the traffic through an fraudulent or compromised DNS server and any web traffic would be zero-rated.

Figure 5



In contrast, the policy depicted in Figure 6 (on the following page) contains additional policy layers that eliminate the ability to trick a charging system via header fraud or DNS spoofing.

By including a DNS check, a fraudster can no longer succeed with committing fraud by routing Internet traffic through a deceptive DNS server, because the deceptive DNS server wouldn't be on the DNS whitelist.

5 <https://www.sandvine.com/resources/whitepapers/application-zero-rating-considerations-and-best-practices-for-revenue-assurance.html>

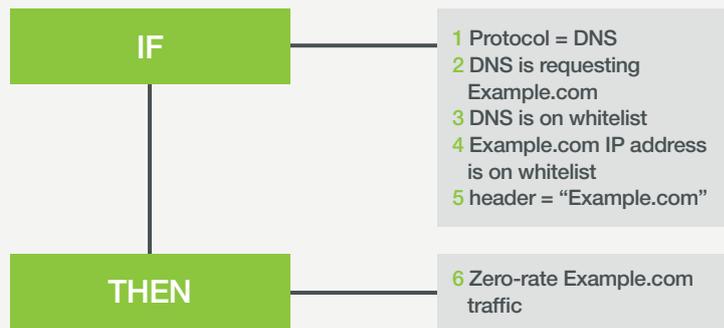


Header fraud, on the other hand, is eliminated due to the IP address whitelisting. Since a user attempting header fraud is deliberately trying to access a website that isn't part of a zero-rated plan, the fraudulent IP address isn't on the IP whitelist.

Clearly, it is important for CSPs to select a vendor with a history of developing successful and fraud-resistant traffic enforcement policies.

Figure 6

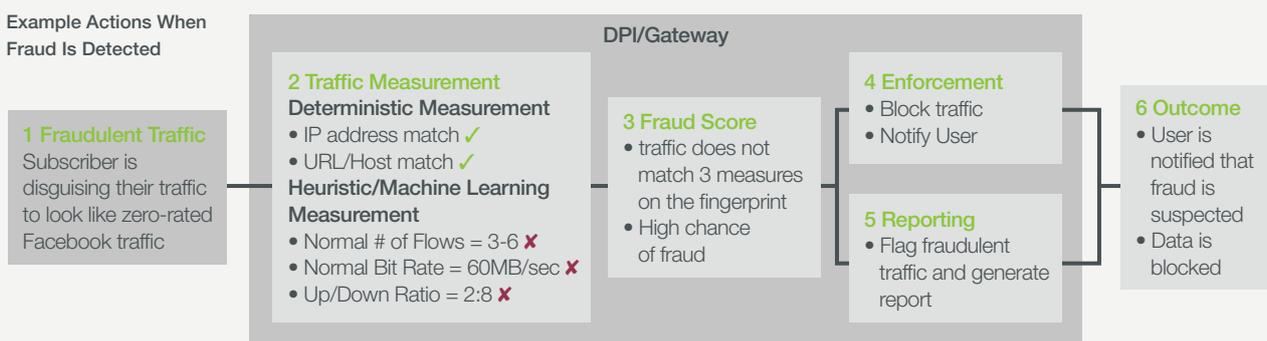
Fraud Resistant Zero-Rating Enforcement Policy



In addition to providing strong zero-rating enforcement instructions, CSPs must rely on the traffic classification profiles to trigger fraud management policies. Depending on the likelihood of fraud (e.g., based on the number and degree of deviations from the application fingerprint), different enforcement actions could be taken automatically by the PCEF or charging gateway. Examples of enforcement actions include: blocking traffic (indefinitely or for a short period of time), notifying the user, and recording the suspected fraud for future analysis. These actions should be configurable for each CSP and should be easily modifiable. Steps 4-6 in **Figure 7** below provide an example of multiple enforcement actions that could be taken when fraudulent activity is identified.

Figure 7

Example Actions When Fraud Is Detected



Reporting Metrics

Accurate reporting metrics are another important aspect of fraud prevention. Determining where fraud is taking place (e.g., the specific zero-rated plans being defrauded) and the pervasiveness of zero-rated fraud (e.g., 5% of all customers) can significantly impact a CSPs fraud prevention strategy. For example, a CSP may choose to completely overhaul zero-rated offerings if fraud is suspected in 30% of subscribers. However, if fraud is equivalent to less than 1% of all subscribers, it may not be worthwhile to adjust traffic enforcement policies.

Additionally, the availability of meaningful metrics means that a CSP's anti-fraud team can analyze and act quickly, which both protects revenue and reduces operational costs.

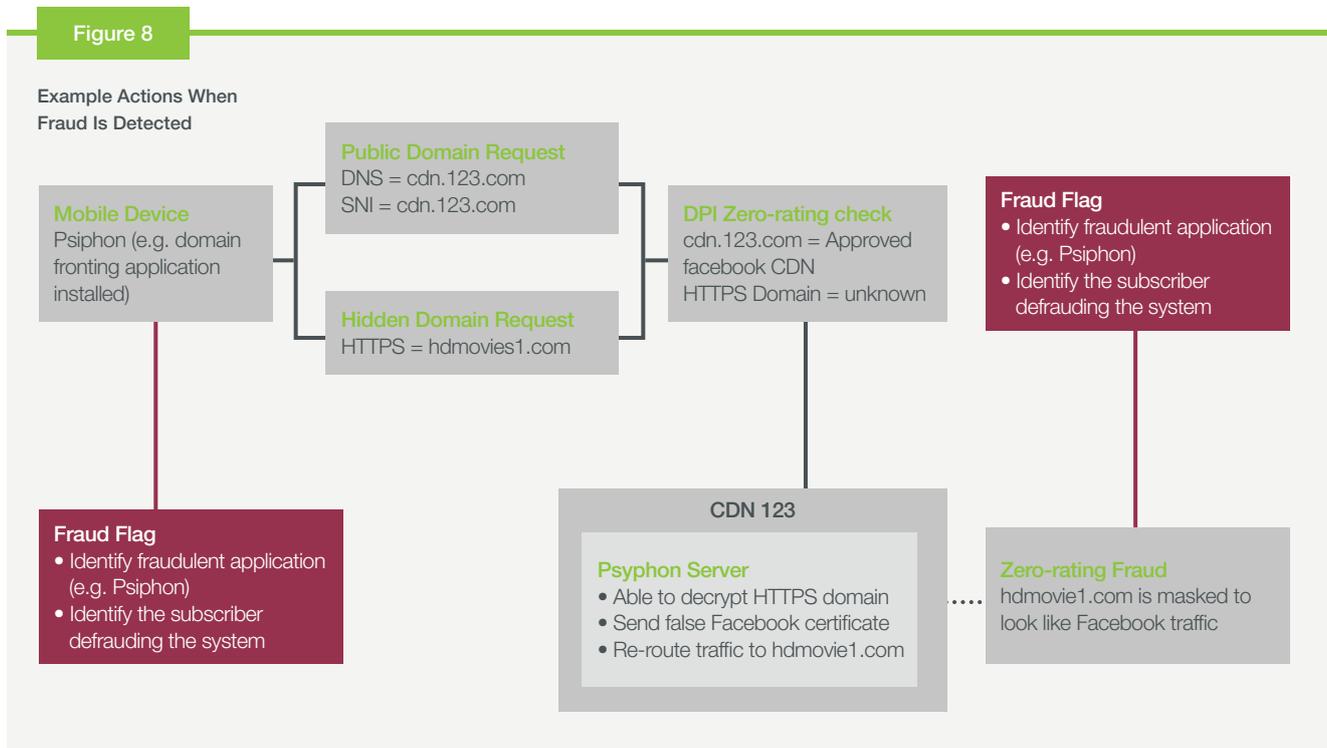


To adequately measure fraud, a CSP must rely on their DPI and/or charging enforcement vendor to both identify and flag fraudulent activity. As noted in the traffic identification section, fraud detection requires advanced traffic recognition techniques like application fingerprinting. Once fraud is detected, it's important that the zero-rated enforcement policy includes a rule to flag all instances of fraud. This reporting flag should enable a CSP to generate a report that includes the following information:

- 1 The application (e.g., Psiphon) used to enable the fraud (if applicable)
- 2 The subscriber(s) using the fraud technique
- 3 The zero-rated application and/or plan being defrauded
- 4 The amount of fraudulent data used

For example, if a fraudster is using Psiphon to mask traffic and defraud a zero-rated Facebook plan, then a CSP should detect the initial Psiphon server connection and flag that as a fraud starting point. Then a CSP could use application fingerprinting to measure how much fraudulent Facebook data was used – see **Figure 8** for a visual representation.

Figure 8



Once a CSP discovers how pervasive the zero-rated fraud is, and how much revenue leakage is occurring, then it's up to them to determine the appropriate corrective action. In our experience, CSPs prefer to quantify the pervasiveness of fraud before enacting preventative measures. Customers typically select this option because it may be more costly to revise zero-rated plans if revenue leakage is miniscule. Due to this preference, we recommend that CSPs analyze fraudulent activity before modifying enforcement rules.

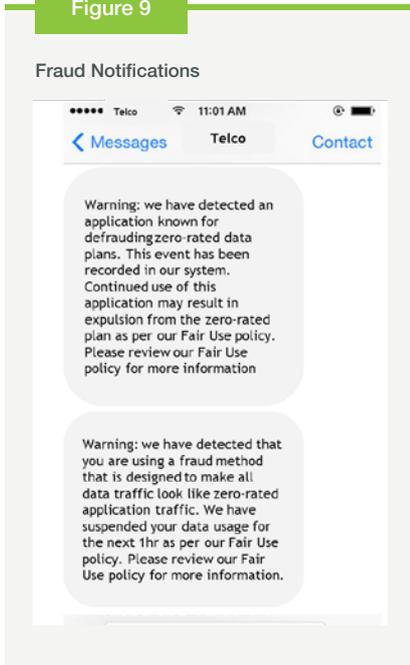
Subscriber Engagement

It's imperative that zero-rated plans are transparent and easily understood by subscribers. Customers need to understand what data is included, and what data isn't included in their zero-rated offering. They also need to understand what activities are permitted, and what activities are not permitted under the zero-rated plan. A full overview on go-to-market messaging for zero-rated offerings can be found in Application Zero-Rating: Considerations and Best Practices for Revenue Assurance⁶.

6 <https://www.sandvine.com/resources/whitepapers/application-zero-rating-considerations-and-best-practices-for-revenue-assurance.html>



Figure 9



To ensure that subscribers fully comprehend what constitutes acceptable zero-rated data usage, we recommend that CSPs post a zero-rating Fair Use Policy that explains:

- 1 Permitted activities (e.g., any zero-rated application receives unlimited data usage for the duration of your prepaid or postpaid plan)
- 2 Prohibited activities (e.g., masquerading data traffic to take advantage of zero-rated offerings is strictly forbidden)
- 3 The consequences of defrauding a zero-rated plan (e.g., if a user is caught defrauding a zero-rated offering by using a masquerading technique or application, then their data service will be suspended for 1 hour; continued fraudulent activity will result in expulsion from the data plan)

We also recommend a progressive discipline approach to fraud, because a subscriber may be unaware that they are defrauding the charging system (e.g., they could be using a masking agent to circumvent a firewall) and should be given a warning before their service is cancelled. When a CSP has identified fraud, notifying the customer is an important next step. This notification should alert the user that fraud has been detected and should explain the corrective action(s). We recommend utilizing a notification system that sends a real-time message to the fraudulent user, using a range of digital channels (e.g., SMS, in-browser overlay, browser push notification, etc.).

By sending a fraud notification in real-time, the user will be more aware of the fraudulent activity that triggered the warning/corrective action. Optimally, the notification should provide a link to the Fair Use Policy so that the customer can better understand how to avoid breaking the terms in the future. Figure 9, below, provides examples of fraud notifications.

CONCLUSIONS

Application zero-rating is a fantastic way for CSPs to offer additional value to their customers and differentiate from the competition. In order to minimize the impact of fraud, CSPs need to adhere to the best practices in **Table 2**.

Table 2

Best Practice	Description
Advanced Traffic Detection Techniques	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • To adequately measure fraud, a CSP must flag and report on all fraudulent activity • Fraud reports should include the following information: <ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • Zero-rated messaging <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • 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



ADDITIONAL RESOURCES

Thank you for taking the time to read this whitepaper. We hope that you found it useful, and that it contributed to a greater understanding of application zero-rating and the challenges and benefits such offerings bring to communications service providers.

In addition to the resources cited in this document, please consider reading these documents related to zero-rating, all of which are available on our website:

- Reasonable Network Management: Best Practices for Network Neutrality⁷
- Online Charging with Diameter Gy: Considerations for Accuracy and Reliability⁸
- Best Practices for Zero-Rating and Sponsored Data Plans under Net Neutrality⁹
- Identifying and Measuring Internet Traffic: Techniques and Considerations¹⁰
- Application Zero-Rating: Considerations and Best Practices for Revenue Assurance¹¹

If you have any feedback at all, then please get in touch with us at whitepapers@sandvine.com.

7 <https://www.sandvine.com/downloads/general/whitepapers/reasonable-network-management.pdf>

8 <https://www.sandvine.com/resources/whitepapers/online-charging.html>

9 <https://www.sandvine.com/resources/whitepapers/best-practices-for-zero-rating-and-sponsored-data-plans-under-net-neutrality.html>

10 <https://www.sandvine.com/downloads/general/whitepapers/identifying-and-measuring-internet-traffic.pdf>

11 <https://www.sandvine.com/resources/whitepapers/application-zero-rating-considerations-and-best-practices-for-revenue-assurance.html>

v20180314

ABOUT SANDVINE

Sandvine helps organizations run world-class networks with Active Network Intelligence, leveraging machine learning analytics and closed-loop automation to identify and adapt to network behavior in real-time. With Sandvine, organizations have the power of a highly automated platform from a single vendor that delivers a deep understanding of their network data to drive faster, better decisions. For more information, visit sandvine.com or follow Sandvine on Twitter at [@Sandvine](https://twitter.com/Sandvine).



USA
47448 Fremont Blvd,
Fremont,
CA 94538,
USA
T. +1 510.230.2777

EUROPE
Birger Svenssons
Väg 28D
432 40 Varberg,
Sweden
T. +46 340.48 38 00

CANADA
408 Albert Street,
Waterloo,
Ontario N2L 3V3,
Canada
T. +1 519.880.2600

ASIA
RMZ Ecoworld,
Building-1, Ground Floor,
East Wing Devarabeesanahalli,
Bellandur, Outer Ring Road,
Bangalore 560103, India
T. +91 80677.43333