

The Core of the Matter: Analyzing Malicious Traffic in Cellular Carriers

Charles Lever
Georgia Institute of Technology
chazlever@gatech.edu

Manos Antonakakis
Damballa
manos@damballa.com

Brad Reaves
Georgia Institute of Technology
brad.reaves@gatech.edu

Patrick Traynor
Georgia Institute of Technology
traynor@cc.gatech.edu

Wenke Lee
Georgia Institute of Technology
wenke@cc.gatech.edu

Abstract

Much of the attention surrounding mobile malware has focused on the in-depth analysis of malicious applications. While bringing the community valuable information about the methods used and data targeted by malware writers, such work has not yet been able to quantify the prevalence with which mobile devices are actually infected. In this paper, we present the first such attempt through a study of the hosting infrastructure used by mobile applications. Using DNS traffic collected over the course of three months from a major US cellular provider as well as a major US non-cellular Internet service provider, we identify the DNS domains looked up by mobile applications, and analyze information related to the Internet hosts pointed to by these domains. We make several important observations. The mobile malware found by the research community thus far appears in a minuscule number of devices in the network: 3,492 out of over 380 million (less than 0.0009%) observed during the course of our analysis. This result lends credence to the argument that, while not perfect, mobile application markets are currently providing adequate security for the majority of mobile device users. Second, we find that users of iOS devices are virtually identically as likely to communicate with known low reputation domains as the owners of other mobile platforms, calling into question the conventional wisdom of one platform demonstrably providing greater security than another. Finally, we observe two malware campaigns from the upper levels of the DNS hierarchy and analyze the lifetimes and network properties of these threats. We also note that one of these campaigns ceases to operate long before the malware associated with it is discovered suggesting that network-based countermeasures may be useful in the identification and mitigation of future threats.

1 Introduction

Malware writers have begun to pay attention to mobile phones. In response, a significant amount of effort has been spent by researchers to characterize malware in mobile applications markets [14, 15, 21, 48, 47]. These efforts have applied a range of static and dynamic analysis techniques on a large number of applications in an attempt to discover malicious code. Market operators including Google and Apple have also invested significant resources in an attempt to prevent malicious applications from being installed on mobile devices and for later removing such applications if necessary. However, for all of these efforts, the extent to which the mobile ecosystem is actually infected by such malware is not well understood. Without such an analysis, it is impossible to determine whether or not current defense mechanisms are having any demonstrable effect.

In this paper, we make the first network level analysis of mobile malware using traffic from a major cellular network. We work from the hypothesis that malicious mobile applications are not different from the bots and malware in the non-cellular world in that they rely on the same core functionality of the Internet in order to achieve scale and robustness. In particular, we began our research with the belief that malicious behavior in the mobile environment similarly relies on the same Internet hosting infrastructure used to support traditional malware activities including propagation and update (e.g., a malware download site), command and control (e.g., communication with infected devices), and data transfer (e.g., a site to upload of stolen data). Should this hypothesis prove to be true, more traditional network-based techniques for detecting and combatting malware [7] can potentially be applied in this new space. As research on botnets and malware have shown, such an Internet-based approach is more effective and scalable than malware-analysis based approaches, particularly in the face of (future) mobile mal-

ware with stealth, obfuscated logic, such as triggered-based behaviors.

We verify our hypothesis experimentally using three weeks of DNS data from a major US-based cellular provider collected over the course of three months. We first show that the vast majority of the hosts resolved in the cellular dataset are also seen in a separate DNS dataset from a non-cellular ISP. After this confirmation, we dig more deeply into the cellular dataset and uncover a number of important results regarding malicious behavior in cellular networks, including the following contributions:

- Known mobile malware samples are virtually unseen:** We extract DNS domains from large public and private datasets of mobile malware and specifically search for their resolution in our dataset. Our analysis demonstrates that only a vanishingly small number of mobile devices appear to be infected: 3,492 out of 380,537,128 devices, or less than 0.0009% of the population. This lends credence to the argument that while the mechanisms market operators implement to protect users from malware may be bypassable [9, 28], malware writers are failing to infect mobile devices with much success. Like any application developer, a malware writer faces the challenging task of developing a popular application (or a malicious application in disguise) that will be downloaded by a large population. That is, the probability that a user will download an unknown malicious application is very small. In addition, a legitimate application market will remove any known malicious application, further reducing the probability of a mobile malware being downloaded.
- Compare traffic to suspicious hosts from iOS against all other devices:** Common opinion argues that the closed nature of Apple’s App Store and operating system make devices in this ecosystem more secure. However, our analysis demonstrates that approximately 8% of iOS-based devices communicate with known suspicious hosts, virtually the identical frequency as all other mobile platforms. Accordingly, users of these devices do not appear to be any more or less likely to communicate with potentially malicious hosting infrastructure than other users. To say the least, iOS does not appear to provide any more effective mechanisms to prevent users from engaging in unsafe activities.
- Observe campaigns with mobile malware clients:** We obtain traffic from the upper layer of the DNS hierarchy and analyze two major threats with mobile malware clients. We see that the lifetime of these two threats lasts on the order of months, and that the criminal operators make use of network agility and move

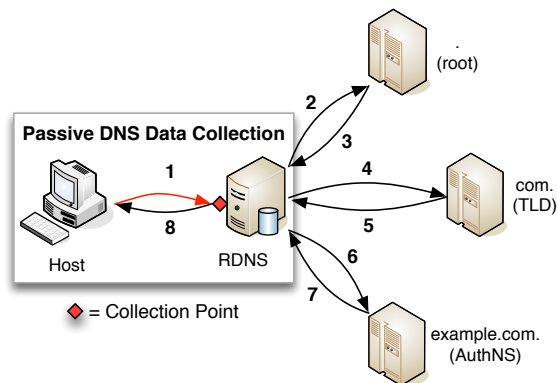


Figure 1: Example of the domain resolution process.

their hosting infrastructure over time (e.g., changing domain names and IPs). Finally, we also note that one of these campaigns cease operating long before the mobile malware associated with each campaign is ever identified, lending credence to the possibility that network-based countermeasures may help identify and mitigate (e.g., via DNS blacklisting) such threats faster than the analysis of mobile malware.

The remainder of the paper is organized as follows. Section 2 provides some background on DNS and related works. Readers familiar with this may skip to Section 3, which describes our dataset collection methodology and the internals of the reputation propagation framework along with the statistical features collected from the cellular DNS data. Section 4 describes the internals of our knowledge base and passive DNS data collected from a large non-cellular US ISP. In Section 5, we present the findings of our measurements and experiments.

2 Background

2.1 DNS

DNS is a backbone protocol for the Internet that maps easy-to-remember domain names to IP network addresses. The domain name space is arranged as a tree, beginning with a root node. In the DNS hierarchy, under the root node are the *top-level domains* (TLDs), and under the TLDs are the *second-level domains*, and so on. Common TLDs include *com.*, *net.*, and *uk.*. A *fully qualified domain name* (FQDN) includes all domain levels that describe the node in the DNS tree; for example, the FQDN *www.example.com.* contains the TLD (*com.*), the 2LD (*example.*), and the 3LD (*www.*). A large portion of the domains are registered directly under a TLD. In some cases, however, this is not possible; therefore, domains under which users can directly

register a new domain name are often considered *effective TLDs*. The canonical example is that many domains in the UK are registered under *co.uk*. and not under *uk*. directly.

The basic type of information link in DNS is the *resource record* (RR). DNS defines a number of RR types. For example, an A-type RR links a domain name with an IPv4 network address, while a CNAME-type RR links a domain name with another “canonical” domain name [31, 32].

RRs are returned in response to a DNS query from a requester. Figure 1 illustrates the DNS query process from a host for the A-type record for *example.com.*. A DNS query is initiated by a DNS *resolver* running on a host. This application is responsible for generating some sequence of queries and translating the responses to arrive at the requested resource. There are two parts in a typical DNS resolution request: the recursive and iterative part. In a typical use, an end system will issue a recursive request using a *stub resolver* to a dedicated *recursive DNS resolver* (RDNS) (Step 1, Figure 1). In a recursive request, the RDNS is charged with completing the iterative portion of the DNS resolution process. It will communicate with the necessary remote name servers (NS) and returns a DNS answer, from the authoritative NS for the requested domain, to the stub resolver in the form of an RR-set. In the case of Figure 1, the RDNS sends iterative requests to the various levels of the DNS hierarchy (Steps 2–7). In Step 7, the RDNS receives the authoritative answer for *example.com.*, and sends it to the requester (or stub resolver) in Step 8, completing the DNS resolution. The RDNS will typically cache the RR locally for up to some period, the Time To Live (TTL), specified in the RR. This improves efficiency and reduces the load on the DNS infrastructure.

2.2 Passive DNS Monitoring

Since a RDNS mediates all requests from a client’s stub resolver, it is possible to perform passive DNS (pDNS) data collection of DNS queries received at the RDNS. This pDNS data collection typically includes all of the information associated with the successfully resolved DNS queries by the RDNS.

There are several benefits to using pDNS monitoring for DNS analysis. Malicious queries are able to be logged and analyzed without alerting the owners of the malicious domains (unlike DNS probing [30, 26]). Another benefit of pDNS data collection is that it can allow the discovery of malicious domains not previously known to exist on DNS blacklists (DNSBL) [7, 8, 10] and does not require previous knowledge of the domain’s existence. A potential drawback of pDNS data collection at the local RDNS level is that the data collected will be limited by the amount of traffic handled by the RDNS. Therefore, it is important to collect a large number of queries from a geographically diverse set

of RDNS servers.

3 Methodology

In this paper, we analyze DNS data generated by devices subscribed to a major US cellular carrier. For simplicity, we use the term “mobile devices” to describe smartphone and tablet platforms (e.g. Android, iOS, and others). Our primary focus is Internet-based *hosts* contacted by mobile devices using DNS. By *host* we mean a remote IP address included in a successful *A-type* DNS domain name resolution. We focus on the hosts as opposed to domain names explicitly because of the extensive forensic evidence already connected with malicious behavior in the Internet (e.g., C&C, drive-by, PPI, etc.). We link the hosts observed in DNS resolutions within a cellular network with historic evidence of Internet abuse and reveal the extent to which mobile devices contact hosts historically associated with known malicious behavior.

Our monitoring point is a single sensor that aggregates traffic collected from several RDNSs providing DNS resolution services for mobile devices located across many different US states. We perform two high-level tasks in order to characterize malicious behavior generated by mobile devices. First, we remove traffic generated by non-mobile devices (i.e., laptop and desktop machines also serviced by this provider via hotspots and “cellular connect cards”) and then attribute each request to a specific mobile device; second, we perform a reputation analysis of the RRs associated with these mobile requests and classify traffic as having benign, malicious or unknown reputation. This workflow is shown in Figure 2, and the details of each task are provided in the remainder of this section.

We use the following notation to more formally describe our workflow. A resource record (RR) refers to either an *A-type* or *CNAME-type* record and its corresponding RDATA [31, 32] tuple of $(qname, ip)$ and $(qname, domain)$ respectively. We represent a DNS query, q_j , for a domain name, d , and the related DNS response, r_j , as the tuple $\mathcal{Q}_j(d) = (T_j, R_j, IPs_j)$, where T_j identifies the epoch in which the query-response was observed, R_j is the unique identification of the device that initiated the query q_j , d is the queried domain, and IPs_j is the set of resolved hosts as reported in the response r_j .

3.1 Mobile Device Identification Process

Accurately identifying individual devices is challenging due to IP address churn and device roaming in a cellular carrier. In order to ensure that we do not include non-mobile devices in our evaluation, we must first be able to attribute traffic to specific devices. We received proprietary data from the carrier that allows us to definitively attribute

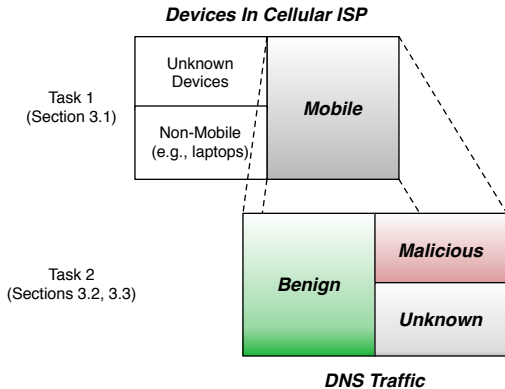


Figure 2: A high-level view of the two tasks required to identify malicious behavior by mobile devices. First, all traffic generated by non-mobile devices is filtered out of our dataset. Second, the remaining traffic is characterized as either benign, malicious or of unknown reputation.

traffic to a mobile device operating on the network for each epoch; however, we cannot correlate devices across epochs.

We formally define the set of mobile devices in an epoch. Let $MDEV_i = \{R_k\}_{k=1..l}$ be a set that contains all l mobile device identifiers (R_k) for epoch i , and $i \in EPOCHS$ ¹. We then define the set of domains resolved by $MDEV_i$ in epoch i as $mDN_i = \{d_k\}_{k=1..n}$, where n is the number of domains (d_k).

Mobile devices are restricted to certain ranges of device IDs that are only associated with operation over a cellular data connection (i.e., 3G or 4G). While this excludes most traditional computing devices, traffic generated by a laptop or desktop using cellular connect cards or tethering through a mobile phone remains in our dataset. We remove this traffic by examining the resource records associated with each device and looking for mobile-specific domains. While there is no official standard, there are some common indicators for mobile-specific domains. One is the use of mobile-specific subdomains such as *m.example.com* or *mobile.example.com*. Another method is the use of the URI path as a mobile indicator, but since we are working with DNS data, we do not have access to this type of indicator. Table 1 provides a subset of mobile indicators.

In addition to mobile indicators, some domains are strongly or exclusively associated with mobile applications. Such sites include mobile ad networks, mobile application programming interfaces, and mobile services. For example, Google’s AdMob [1] advertising network is only supported on Android and iOS, so devices that contact the AdMob

¹We define the set of epochs as *single days* falling within the following ranges $EPOCHS = \{4/15/12 - 4/21/12, 5/13/12 - 5/19/12, 6/17/12 - 6/23/12\}$

Domain name	Indicator Type
<i>m.example.com</i>	Subdomain
<i>mobile.example.com</i>	Subdomain
<i>android.example.com</i>	Subdomain
<i>iphone.example.com</i>	Subdomain
<i>ipad.example.com</i>	Subdomain
<i>touch.example.com</i>	Subdomain

Table 1: Examples of popular mobile indicators.

network are almost certainly mobile. Certain mobile APIs use unique domains that are easily identifiable. Apple’s Push Notifications, for instance, use a set of mobile domains (e.g., **.courier.push.apple.com*) reserved specifically for the push notifications service. Additionally, services like HeyTell [2] provide push-to-talk functionality for mobile devices via a mobile application. Devices communicating with these types of services should be almost exclusively mobile.

We first exclude all devices R_k that query any of six domain names d in an epoch j related to standard operations of Windows operating systems². Devices are labeled as mobile only if they contact a domain d with a mobile indicator or domains strongly associated with mobile applications or services. We are able to identify 132,516 unique domain names that fit this former category.

Through the combination of filtering based on device IDs and mobile domain inference, we are able to identify devices as either mobile or non-mobile with high confidence. In the small number of cases where overlap exists, we tag a device as unknown and do not consider its behavior as reflective of mobile devices. By conservatively choosing devices in this manner, we strongly reduce the likelihood of selecting traditional computing devices that are connected via mobile broadband cards or tethering. However, if a traditional computing device is visiting a mobile resource, we would falsely label that device as mobile. We note that this scenario is exceedingly rare as browsers generally direct users to the appropriate version of a website.

3.2 Filtering Benign Domains

At this point in our workflow, we can label a device and associate it with RRs. We can then begin the second task in Figure 2: classifying each of the resource records requested by mobile devices. The first step in this process is identifying and removing known benign traffic from our dataset. Our methodology again remains **conservative**, aggressively reducing false positives potentially at the cost of increasing false negatives. We achieve this by whitelisting all re-

²Two of the most frequently hit domain names in this list are the *time.windows.com* and *download.windowsupdate.com*.

INPUT : $MDEV_j, mDrdc_j$ and all $\mathcal{Q}_j(d_i)$ for every mobile domain name ($d_i \in mDrdc_j$) observed in epoch j .

Let $HOSTS_j = \emptyset$, be the set that will contain the unique hosts (or IP addresses) that have been mapped with the domain names in $mDrdc_j$ after the completion of the process.

[1]: $\forall d \in mDrdc_j$:

[2]: Let IPs_j be the set of IPs in the tuple $\mathcal{Q}_j(d)$, if $R_j \in MDEV_j$

[3]: $HOSTS_j \cup IPs_j$

OUTPUT: $HOSTS_j$

Algorithm 1: The algorithm used to obtain the set of IP addresses (or hosts) that represent the hosting infrastructure that facilitated resolution of domain names from mobile devices in epoch j .

quests made to the top 750,000 effective second level domains (e2LDs) according to Alexa [6]. However, we must note that we do not whitelist domains associated with dynamic DNS (DDNS) providers given their common use by network malware. Intuitively, such broad whitelisting removes the most popular sites (and e2LDs) as they are more likely to be trustworthy and less likely to be intentionally malicious. This approach is commonly used in DNS-based reputation and classification systems [7, 8].

We want to remain as conservative as possible and reduce any potential false positives from our dataset. To that end, we further filter benign traffic from our dataset by removing a number of the most popular remaining e2LDs. We compile a list of approximately 800 e2LDs based on the lookup volume of the queried domains. We manually inspect all of them and we classify them as benign. We should note that the lookup volume distribution for the e2LDS follows a power law.

The end result of the whitelist filtering process is a reduced set of domain names $mDrdc_j = \{d_k\}_{k=1\dots n}$, where $mDrdc_j$ is a set that contains all n **not whitelisted** domain names (d_k) resolved by mobile devices in epoch j . In this set we will have domain names in the “Malicious” and “Unknown” categories of Figure 2.

3.3 Feature Extraction

The remaining entries in our dataset now belong to mobile devices communicating with Internet-based hosts with either malicious or unknown reputations. We now describe the features that we extract from these remaining domains, which will allow us to analyze the **hosting infrastructure** supporting these domains. We use Algorithm 1 to find all unique IP addresses (hosts) for all domains in the set $mDrdc_j$ for epoch j , resulting in the set $HOSTS_j = \{IP_k\}_{k=1\dots n}$. For ease of understanding, the entire feature extraction process is summarized in Figure 3.

We compare the traffic observed in the cellular carrier (only from the mobile devices) against a pDNS data collection from a non-cellular ISP. Let $f_{pdns}(d) = \{ip_k\}_{k=1\dots n}$ be a mapping function that takes a domain name d as input and returns a set of routable IP addresses that have been historically linked with d .

The $f_{cell}^j()$ function returns passive DNS data from the DNS traffic in the cellular carrier over an epoch j . Let $f_{cell}^j(ip) = \{d_k\}_{k=1\dots n}$ be a function that receives an IP address ip as input and returns a set of related historic domain names (d_k) observed in the cellular network during epoch j from mobile devices in set $MDEV_j$.

The function $f_{mal}(ip) = \alpha$ returns the number α of unique malware samples that IP address ip has been associated with over the past 19 months. The association could be direct (i.e., the malware contacts the IP address/host) or indirect (i.e., the malware looks up a domain name that resolves to that IP address/host) as shown in Figure 3. We similarly define $f_{url}(ip) = \beta$ and $f_{pbl}(ip) = \gamma$, as the functions that return the number of malicious URLs (β) and malicious entries (γ) in public blacklists.

For every host (ip) in the set $HOSTS_j$ we extract the two groups of features: passive DNS and evidence features. *At the end of the feature extraction process, we obtain statistical historic passive DNS and malicious evidence-based observations (used in Section 5.1) for the set of hosts in $HOSTS_j$.*

Passive DNS Features (PF) We collect two features from this group. They are simply the number of elements in the sets $f_{pdns}(ip)$ (i.e., related historic non-cellular domains) and $f_{cell}^j(ip)$ (i.e., related historic cellular domains) for an address ip .

Evidence Features (EF) We compute a total of three features from this group. These features describe direct reputations of the IP addresses in the set $HOSTS_j$ (during the epoch j). We compute three features for each IP address: (i) $f_{mal}(ip)$, the count of unique malware associated with ip , (ii) $f_{url}(ip)$, the count of URLs associated with ip , and (iii) $f_{pbl}(ip)$, the count of public blacklisting incidents associated with ip .

Both PF and EF feature families represent the basic building blocks of DNS reputation systems [7, 10]. We select them to understand (i) the extent that malicious hosts currently serve mobile-related DNS resolutions and (ii) the extent that the infrastructure used to resolve mobile-related domain names is already present in passive DNS data collections from non-cellular networks. In particular, the PF feature family, which is based on passive DNS data, will show to what extent the hosts from mobile RRs (directly or indirectly) can be associated with DNS resolutions from a

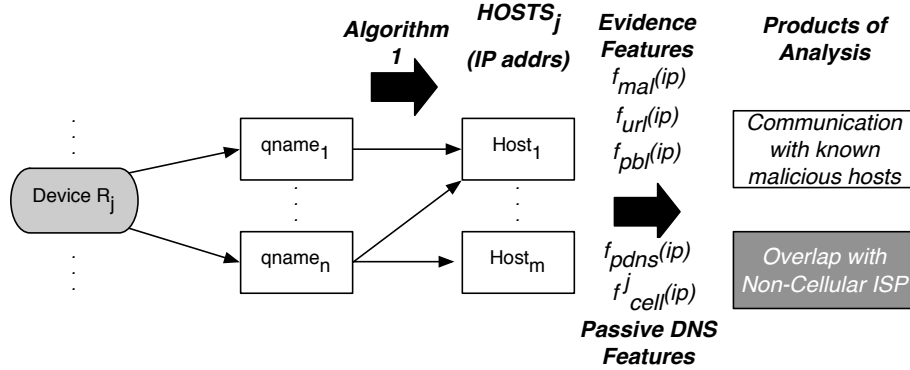


Figure 3: Determining the communication patterns for each mobile device (R_j). Each qname requested by R_j is converted into an IP address via Algorithm 1. This list of IP addresses ($HOSTS_j$) is then processed for Passive DNS Features (PF) to determine overlap with traffic from our non-cellular ISP and for Evidence Features (EF) to determine the presence of communications with known malicious domains.

non-cellular ISP. The EF features, which are based on historic reputation information, will show us to what extent the already tainted Internet hosting infrastructure is currently used directly by mobile devices. Additionally, we perform one more level of filtering in which we evaluate the malicious hosts identified by the EF feature family using the Notos [7] reputation system; we remove any hosts identified using the EF features if Notos does not produce a reputation score below our chosen threshold.

4 Dataset Summary

This section describes the datasets used in our analysis. These include pDNS data collected from a major US cellular carrier, pDNS data collected from a major US non-cellular ISP, and a database of malicious evidence built from several classes of malicious information.

4.1 DNS

Cellular We observed DNS traffic from a cellular data network on twenty-one days over a three month period. This data was passively collected from a single sensor that aggregates information from several cities.

4.1.1 Observations from the Cellular Carrier Traffic

Table 2 provides insight into the number of *unique* RRs, domains, and hosts seen over the twenty-one (single day) epochs. For each record type, there are two columns that specify the total number of unique records seen for the given week and the number of new records not seen in any prior week. Intuitively, the number of new records seen should decrease over time, and Table 2 shows that the influx of RRs, domains, and hosts does follow this pattern.

Non-cellular The non-cellular pDNS data was collected from seven different sensors located across the US over more than 15 months. Due to the extended collection period, this dataset presents a substantial volume of traffic that can be used to provide historical context for domains and hosts of interest. In particular, we can use this data to make inferences about the hosting infrastructure of a particular domain or tie specific hosts to their related domains.

4.2 Devices

Devices seen in the cellular dataset accessed the network via a cellular data connection. Consequently, these devices should fall into three general categories: smartphones, tablets, or mobile broadband devices. The first two categories include devices such as Android and iOS phones and tablets. A mobile broadband device includes any device accessing the network via a mobile broadband card or tethering to another device’s cellular data connection. This could include traditional computing devices such as desktop or laptop computers.

As discussed in Section 3.1, we are conservative in the classification of mobile devices; a comparison of the total devices and what was classified as a mobile device can be seen in Table 2. Most importantly, this table shows that our estimate of mobile devices is conservative; we classify only 79% of devices seen as mobile.

4.3 Evidence

We analyze cellular DNS traffic with an evidence database composed of three general classes of non-mobile malicious evidence: public blacklist data (PBL), phishing and drive-by download evidence (URL), and hosts accessed by known malicious applications (MAL). In addition to the

	Duration (hours)	RRs		Domains		Hosts		Devices	
		Total	New	Total	New	Total	New	Total	Mobile
4/15-4/21	168	8,553,155	8,553,155	8,040,141	8,040,141	2,070,189	2,070,189	157,286,931	121,497,066
5/13-5/19	168	9,240,372	4,498,765	8,711,704	4,042,009	2,168,266	606,467	169,561,760	136,292,358
6/17-6/23	168	8,660,555	3,246,194	8,109,536	2,745,999	2,050,168	377,048	153,525,716	122,747,704
Total	504	26,454,082	16,298,114	24,861,381	14,828,149	6,288,623	3,053,704	480,374,407	380,537,128

Table 2: Listing of unique RRs, domains, hosts, and devices seen in cellular dataset.

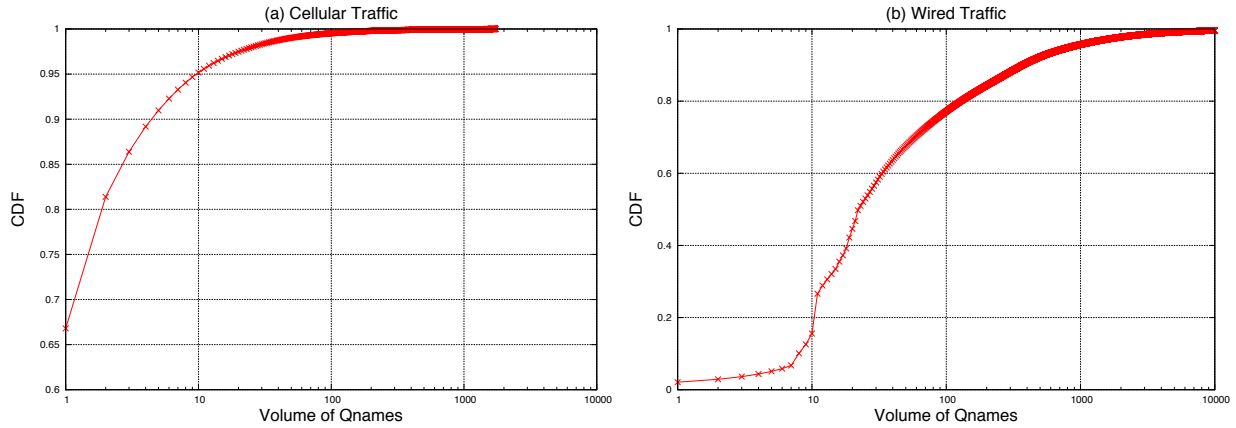


Figure 4: Per host (in $HOSTS_{all}$) qname distribution of DNS evidence. These experiments demonstrates that the hosts in $HOSTS_{all}$ have a significant historic presence in pDNS data collected from non-cellular networks.

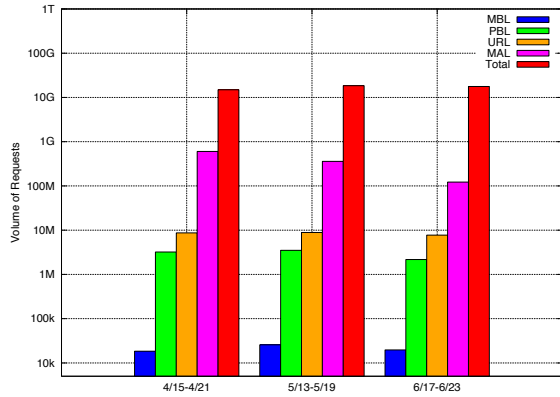


Figure 5: Volume of requests to domains with malicious evidence visited by mobile devices in cellular network.

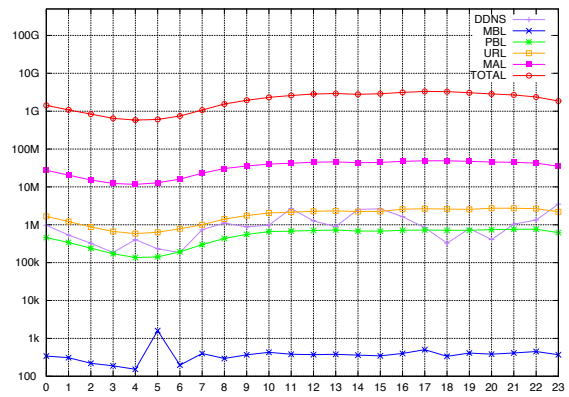


Figure 6: Hourly analysis of request volume for various types of domains observed from mobile devices.

non-mobile evidence database, we also used a mobile blacklist (MBL) containing 2,914 domains known to be associated with mobile malware or mobile malware operators. Figure 5 shows the volume of DNS lookup requests from mobile devices in our cellular dataset that could be directly linked with different classes of non-mobile and mobile evidence.

Additionally, a diurnal analysis of the volume of requests that can be directly associated with different types of mali-

cious evidence can be seen in Figure 6. It is important note that our cellular DNS traffic sensor aggregates data from several different locations in different time zones. Therefore, patterns may be less pronounced than if all data was collected from a single location. Looking at the total volume of requests for each hour, Figure 6 shows that the volume of traffic starts to increase between hours five and six and begins to gradually decrease between hours seventeen and eighteen. These hours approximately correspond

to people waking up in the morning and traditional dinner hours in the evening. With the exception of DDNS traffic and a spike at hour five for MBL requests, the other classes of domains appear to follow the total hourly volume pattern.

5 Results

We now present the results of our experimental evaluation. We begin by analyzing the traffic observed in the cellular carrier, first characterizing the data from the cellular carrier in isolation, and then comparing the request patterns to those observed historically in our non-cellular ISP dataset. We then examine the extent to which the hosts observed in mobile resolution requests are directly or indirectly tainted by reputation information collected in non-cellular ISPs.

We continue in Section 5.2 and further focus on evidence of mobile-specific malware. We continue our analysis of directly tainted hosts in mobile resolutions by determining from which mobile platforms those queries originate. Then, we examine all of the cellular network queries for mobile-malware specific domains to determine the extent of the presence of mobile malware in the cellular network.

We conclude our results with long term analysis of two known mobile threats in Section 5.3. In particular, we study those two threats from their rise until they become almost completely inactive. Furthermore, we provide the global infection perspective of those two mobile threats. Finally, we examine their hosting infrastructure and how it changes over time.

5.1 Analysis of the Reputation Datasets

We analyze the DNS traffic generated from approximately 380 million mobile IDs over the course of our observation period (Table 2). Note that while we can identify devices consistently within an epoch (i.e., a single day), we can not link devices across days, meaning that this total does not represent the number of unique devices serviced by the carrier. We are specifically interested in unique Internet-based hosts requested by mobile devices. The filtering process results in the set $HOSTS_{all}$ consisting of 2,902,071 unique hosts having the following characteristics: (i) at least one resolution request for each host was observed by the cellular carrier, (ii) the observed DNS requests came strictly from devices classified as mobile, (iii) the hosts are not associated with any known benign infrastructure, and (iv) the host is a routable IP address. We obtain passive DNS information on each of these hosts from our historical non-cellular ISP dataset (via the function $f_{pdns}()$ defined in Section 3.3).

5.1.1 Observations from the Cellular Carrier Traffic

We use the hosts in $HOSTS_{all}$ to perform an in depth examination of their DNS properties in our pDNS data. We do this for both the cellular ISP traffic (Figure 4(a)) and, using a projection into our passive DNS data collection, for the non-cellular ISP traffic (Figure 4(b)).

The hosts in $HOSTS_{all}$ reflect a portion of the hosting infrastructure that support unknown or malicious types of DNS resolutions in the mobile carrier. We project this set into the non-cellular data collection and obtain non-cellular passive DNS data for the hosts in $HOSTS_{all}$. Only 36,338 (or 1.3%) of hosts in $HOSTS_{all}$ are outside the non-cellular passive DNS evidence we have.

Looking a bit closer, Figure 4(a) shows the distribution of unique hosts in the set $HOSTS_{all}$. We see that more than 18% of the hosts requested by mobile devices are associated with only a single domain. Furthermore, Figure 4(b) shows that 98.7% of hosts in the set $HOSTS_{all}$ have at least one historically associated domain name (according to the passive DNS data collection from the non-cellular ISP). This simply means a sufficiently large non-cellular pDNS data collection can be used to amplify the DNS information for hosts observed in DNS resolutions from non-cellular networks.

If we assign Notos [7] reputation scores to the RRs that have IPs in the $HOSTS_{all}$ set, and use reputation threshold of 0.8 (or above 80% probability of the RR being malicious), we obtain 51,503 domain names (3,636 distinct IPs) as likely suspicious. Only 18 domain names (13 unique hosts) of them have been listed in mobile black lists until the day of the submission. While 0.8 is a conservative threshold, these results could be used as an indicator that the malicious hosting infrastructure observed in cellular networks is already present in reputation and passive DNS observations from non-cellular networks.

These findings are already valuable. Given the significant overlap between the domains requested by devices in cellular and non-cellular providers, and the historical information regarding the reputation of the hosts in $HOSTS_{all}$ (as discussed in this Section and Section 4.3), we can conclude that the DNS infrastructure (malicious or not) is being reused in cellular networks. Moreover, the scores assigned to hosts by DNS reputation systems can potentially serve as filtering functions for applications when they are submitted to mobile markets.

5.1.2 Observations from Mobile Markets and Malware Datasets

We now characterize application markets and datasets of known mobile malicious applications. Specifically, we examine all of the applications in Proandroid, So-fandroid, Anzhi, Ndoo and the top 500 free applica-

Market/Dataset Name	Market Country	Date of Snapshots	#Unique Apps	#Unique domains	#Unique IPs
Google Play*	US	9/20/11 and 1/20/12	26,332	27,581	47,144
SoftAndroid	RU	2/7/12	3,626	3,028	8,868
ProAndroid	CN	2/2/12, 3/11/12	2,407	2,712	8,458
Anzhi	CN	1/31/12	28,760	11,719	24,032
Ndoo	CN	10/25/12, 2/3/12, 3/6/12	7,914	5,939	14,174
Contagio	—	3/27/12	338	246	2,324
Zhou et al.	—	2/2012	596	281	2,413
M1	—	3/26/2012	1,485	839	5,540

*Top 500 free applications per category only

Table 3: M&A Datasets

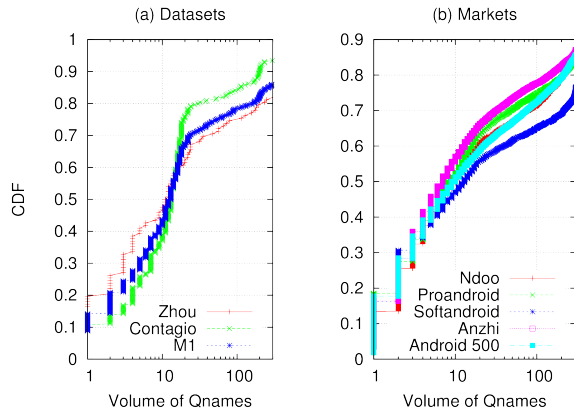


Figure 7: Qname distribution per passive DNS (from non-Cellular networks) evidence gathered from the $m\&a$ dataset. In spite of their geographic diversity, the requested qnames in all of these subsets follow a similar distribution.

tions from each category in Google Play. We also analyze three datasets containing known mobile malware, including all of the malware samples from the Contagio blog, 596 apps from the malware dataset described in Zhou et al. [47] and a third dataset which we refer to as $M1$ that was provided confidentially by an independent security company. $M1$ contains both highly suspicious and confirmed malicious applications. We refer to this collection of datasets as the malware and applications ($m\&a$) datasets ($\{Android500, Proandroid, Softandroid, Anzhi, Zhou, Contagio, M1\}$). We statically extract domain names from application code in the $m\&a$ dataset to create a set of domain names we call $DN_{m\&a}$. Using the $f_{pdns}()$ function, we record the unique hosts historically associated with each entry in $DN_{m\&a}$ from our non-cellular pDNS dataset. Table 3 shows the breakdown of each of these datasets.

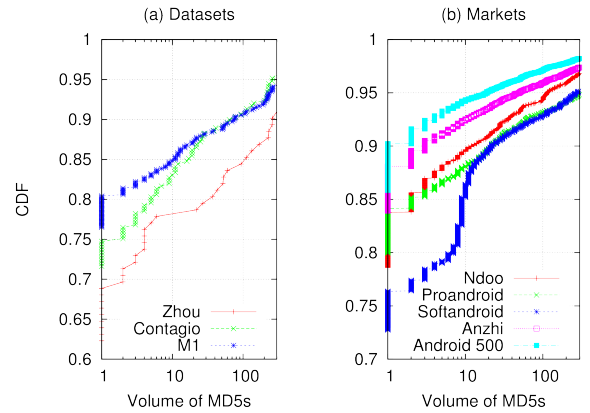


Figure 8: Volume of malicious MD5 evidence associated with unique qnames seen in the $m\&a$ dataset.

Figures 7(a) and (b) show the distribution of qnames for each of the subsets of the $m\&a$ dataset. We observe that at least 90% of the hosts in $HOSTS_{m\&a}$ are present in our non-cellular pDNS dataset, and in some cases (see Figure 7(b)) all of them are present. Despite the geographical diversity of the markets we examine, in Figures 7(a) and (b) we see that all $m\&a$ datasets have very similar distributions of qnames per host in $HOSTS_{m\&a}$. Furthermore, more than 50% of these hosts have at least seven domain names historically associated with the (non-cellular) pDNS dataset.

Figure 8 shows the direct relationship to hosts historically associated with malware for each of the subsets of the M&A datasets. We observe that the 38% of the hosts retrieved from the $HOSTS_{Zhou}$ dataset, which corresponds to applications from Zhou et al. [47] dataset, have been historically associated with more than one malware samples (or MD5). Additionally, approximately 23% of hosts in $HOSTS_{softandroid}$ dataset, can also be linked with more than one malware samples.

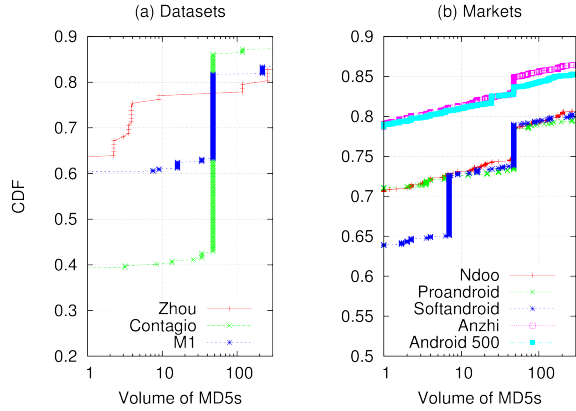


Figure 9: Volume of malicious MD5 evidence associated with unique qnames in the $m\&a$ dataset after projection through non-cellular pDNS data collection.

Somewhat unsurprisingly, Google Play (see Figure 8(b), Android 500 class) has the lowest percentage of applications contacting hosts historically associated with malware. Only a 10% of the domain names observed in the class Android 500 in Figure 8(b), have more than one malware sample associated with the host the domain names resolved to historically. These numbers increase in Figure 9, which shows the projection of $HOSTS_{m\&a}$ through our passive DNS data collection from non-cellular networks using $f_{mal}()$. As previously mentioned, such indirect results are prone to false positives due to phenomena like parking IPs and sinkholes; however, the inclusion of such hosts in an application could easily serve to trigger analysis by market operators.

In summary, mobile applications observed from a variety of marketplaces and malware datasets use much of the same Internet infrastructure as the non-cellular DNS resolutions. This observation is similar to the analysis we made in Section 5.1.1, when we examined DNS traffic from a large cellular provider. Perhaps the most important observation from the $m\&a$ datasets is that some market-places (e.g., Softandroid) contain mobile applications that tend to be more directly and indirectly related with known malware-tainted hosting infrastructure.

5.2 Mobile-malicious activity in mobile networks

This subsection discusses two distinct but related phenomenon: 1) the relationship between mobile platforms and requests for tainted hosts, and 2) the presence of queries for domains facilitating malware that targets mobile platforms.

Device Platform	% Total requests by mobile device	% Population requesting tainted hosts	% Total tainted host requests
iOS	31.6%	8.8%	33.2%
All other mobile (Android, etc.)	68.4%	8.2%	66.8%

Table 4: Tainted Hosts and Platforms

5.2.1 Tainted hosts requested by mobile platforms

Table 4 presents a breakdown of which platforms correspond to what proportion of total mobile device population, what proportion of the device population requested tainted hosts, and which platforms are responsible for tainted host requests. This data is presented for iOS devices and other mobile devices such as Android and other indistinguishable platforms (e.g., potentially Symbian, Windows Mobile, unverifiable iOS devices). We separate iOS devices from the rest because they are easy to reliably identify by searching for push notification domains.

The first column in Table 4 shows the contribution of each platform to the total population of mobile devices. The majority of devices were unidentifiable, but roughly one third of the devices could be identified as iOS. The second column shows the percentage of devices for each platform that requested domains that point to tainted hosts. We observed roughly 8% of iOS and 8% of other mobile devices issued at least one request that pointed to a tainted host.

Finally, the third column shows how platforms contribute to the overall number of tainted host requests. It is interesting that each class of devices contributed proportionally to the number of total tainted host requests. This data shows that, from the network perspective, iOS devices reach out to similar numbers of tainted hosts as other devices.

Figure 5 provides an overview of different threat classes that are present in our PBL evidence set. Each of the different classes shown in this table were seen in our cellular pDNS dataset and was visited by a mobile device. It is interesting that this list includes several prominent desktop malware classes such as Zeus and SpyEye in our filtered mobile traffic.

5.2.2 Mobile malware in the cellular network

To answer the question "How prevalent is the threat of malicious mobile applications in the cellular network?", we scanned all (not mobile-device only) cellular network DNS

Threat Class	# Associated Hosts
Artro	1
Backdoor.Tofsee	1
DMSSpammer	2
FakeAV	1
MalwareDomainList	1386
MalwarePatrol	203
Misused	69
Phishing	383
SC	4
SpyEye	183
Worm.Palevo	30
Zeus	1083

Table 5: Threat Class for Tainted Hosts

data with a blacklist of 2,932 domains known to be associated with mobile-malware or mobile-malware operators. We then focused on those domains that are known to have directly facilitated mobile malicious activity, not merely associated with it. Examples of this malicious activity include distributing malicious applications, exfiltrating sensitive data without user consent, and command and control services.

We focused on 19 unique domains present in our cellular pDNS data. These malicious domains are associated with 10 unique malware families; all of these are Android applications. 9 of these 10 malware families were publicly disclosed *before* any of our epochs — meaning they were still queried after they were known to be malicious by security researchers, antivirus companies, and market providers. [34, 33, 46, 11, 38, 3, 4, 5, 39]

Table 6 shows the mobile malware families with domains seen in the cellular network, the number of domains known to facilitate the malicious activities of those families, the number of devices of any type and the number of mobile devices that contacted a domain facilitating mobile malware. Note that this data is aggregate across days and all domains facilitating mobile malware; we cannot identify the same device across epochs, so this is an upper bound of devices that contact a domain. The effects of our conservative mobile classification process are apparent — only a fraction of devices that we classify as mobile contacted any domain that facilitates mobile malware compared to all devices.

The most prevalent malicious family in the network was FakeDoc. This is a potentially unwanted application (adware) that steals a user’s Google account and other potentially sensitive information. FakeDoc was discovered in the Android market on October 19, 2011 well before our traffic epochs. Despite being flagged by several antivirus products [4], 5,417 devices contacted the domain used by FakeDoc for malicious activity.

The second most popular malware family was NotCompatible. NotCompatible is a trojan application that acts as an open network proxy. Unlike the other families in Table 6, NotCompatible is spread through compromised web pages with hidden iFrames that point to a download site for the app. NotCompatible was disclosed on May 2, 2012 [46].

Even considering an upper bound, the overall traffic to domains associated with mobile malware is quite small. Only 9,033 devices of any type out of a total of over 480M million (0.001%), and 3,492 devices out of a total of 380M confirmed mobile devices (0.0009%) contacted a domain known to facilitate mobile malware. The top two threats present in our data present a stark contrast in functionality, time of known activity, and method of distribution. Despite these differences, neither presented significant activity levels during our measurement epochs.

A number of insights can be gleaned from this data. First and foremost, mobile malware is a real threat to users in the United States, despite the fact that malware researchers find many of their samples in non-US markets. Even though the threat is real, it is minimal. It is important to note that the overall size of all infected populations indicates that mobile malware is far from reaching the scope or severity of desktop malware. This may be attributed to moderated markets, security architectures of mobile platforms, and the relative lack of opportunity an infected device can provide a malware author.

The low volumes of traffic from malware distributed through the Google Market indicate that market-based kill switches can be effective at controlling the malware population. However, the *relative* success of NotCompatible calls into question whether the market-based kill switches will be able to control the spread of malware in the future if malware authors eschew markets in favor of other distribution means. Even when markets are used as a distribution channel, mobile malware can be seen in the network long after discovery by researchers and its removal from markets. This finding implies that neither markets nor security products like antivirus tools are able to guarantee a malware-free platform.

5.3 Lifecycle of Mobile Threats

In this section we examine the evolution of two malicious mobile applications, threat ϵ and threat β . With historic data from a large authoritative DNS server, we provide insights on how the threats developed over time as well as geographical properties of the requesters requesting domains associated with these threats. We conclude by providing some insight on the hosting infrastructure these mobile threats used throughout their lifetime, especially at the peak of their activity.

Malware Family	# Assoc. Domains	#Devices (Any type)	#Devices (Mobile only)
DroidDreamLight*†	3	150	44
DroidKungFu*	1	19	6
FakeDoc*†	1	5417	2145
Fatakr*	1	328	151
GGTracker*	3	1	1
Gone60*†	1	1	1
NotCompatible	3	2198	762
Plankton*†	4	686	286
Malware β *	1	18	1
WalkInWat*	1	215	95

* Disclosed before any of our epochs

† Distributed in Google Play market

Table 6: Malicious Apps with Domains in Mobile Network

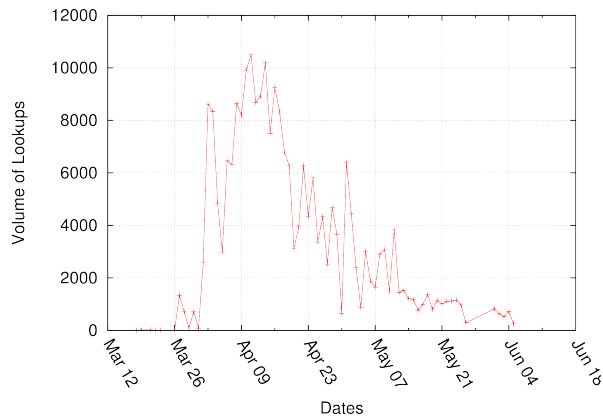


Figure 10: DNS request volume for threat ϵ (2011)

Threat ϵ is an Android application that masquerades as a legitimate client to a popular Internet streaming media service. When run, the application presents a credible login screen. When the user attempts to log in, the application displays an error message and closes. In the meantime, it has sent the user’s credentials to domain $qname_\epsilon$ in an HTTP request. This threat was publicly disclosed by a major anti-virus company in October 2011.

Threat β is an Android application that starts a service after reboot that periodically contacts a C&C server hosted on domain $qname_\beta$. The service will respond to commands received from the C&C or via SMS. One command causes the application to sign all contacts up to an on-line mailing list, while another command has the application send infected download links to all contacts via SMS. These links are on a different domain than $qname_\beta$. The application will automatically respond to received SMS with an offen-

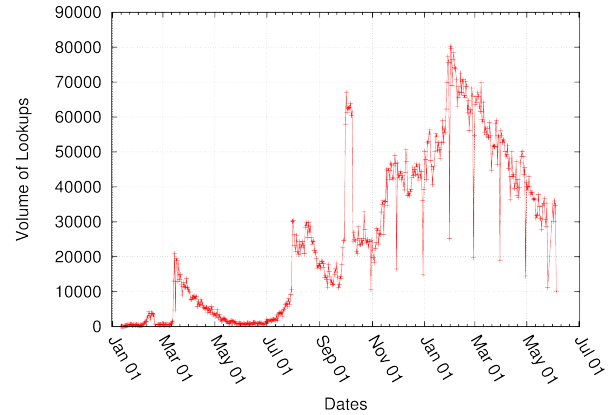


Figure 11: DNS request volume for threat β (2010 to 2011)

sive message, and in certain cases will send offensive SMS messages to all contacts. This threat was publicly disclosed by a major anti-virus company in May 2011.

5.3.1 Lifetime and Infection Scale

Figure 10 shows the daily lookup volumes for $qname_\epsilon$, which acts as a proxy for the victims of threat ϵ . These lookups could be recursive DNS servers, so we cannot make any claims about the size of the overall infected population. The threat was most active on April 12th, but soon after rapidly declines. The first lookup for $qname_\epsilon$ was recorded on March 3rd, 2011, and by June 5th ³ there were DNS requests from 2,731 unique requesters. Table 7 shows the query volume, AS, and country code of the top ten networks that sent requests to $qname_\epsilon$; the majority of these are based in the US. Of note is that this threat seems to have ended well before it was publicly disclosed in October 2011; at the time of disclosure, $qname_\epsilon$ no longer resolved to a routable address.

Figure 11 shows the lifetime of threat β in terms of query volume. This threat became active in January of 2010, and at its peak in February – March 2011 it averaged more than 70,000 DNS requests per day. Over the 14 months that this threat was active, 13,094 unique IP addresses queried the domain name $qname_\beta$. As before, this number cannot be considered an absolute population estimate. Table 7 shows the distribution of the infected populations for mobile threats β and ϵ . We see that a significant portion of the infected population resides in Asia-based networks. We also note that Google (AS 15169) has a heavy impact on the numbers in Table 7 (most likely due to crawling). Threat ϵ was disclosed well past its peak in DNS requests.

³ We have no data from the authoritative DNS server after this date, so we have no visibility into later activity

Threat ϵ			Threat β		
Volume	AS	CC	Volume	AS	CC
816	3356	US	7315	3356	US
112	15169	US	470	3462	TW
97	7132	US	266	15169	US
92	9299	PH	222	4766	KR
67	7843	US	210	7132	US
52	20115	US	160	9299	PH
47	6389	US	139	6389	US
44	7643	VN	127	9121	TR
38	22773	US	122	20115	US
33	24560	IN	115	24560	IN

Table 7: Requester information with respect to autonomous system (AS), country code (CC), and count of unique IPs in the AS (volume).

Threat ϵ			Threat β		
Volume	AS	CC	Volume	AS	CC
11	6389	US	237	6389	US
3	20115	US	28	49544	NL
1	7132	US	28	27589	US
1	13674	US	15	29550	GB

Table 8: Information on the hosting infrastructure used by the two mobile threats.

5.3.2 Hosting Infrastructure

Here we describe in detail the Internet infrastructure used by both threats. Table 8 shows the autonomous systems, country codes, and the number of hosts within each AS that $qname_\epsilon$ or $qname_\beta$ ⁴ pointed to throughout their lifetime. Figure 12 shows how the host pointed to by $qname_\epsilon$ changed over time. The host was primarily located in AS 6389, but for brief periods of time the domain resolves to hosts outside AS 6389. Comparing the activity of threat ϵ (Figure 10) to the changes in the host (Figure 12) reveal that host changes were correlated with activity peaks (as seen in April 2011). Figure 13, shows host changes over-time for threat β . Like threat ϵ , the host infrastructure was relatively stable until peak activity (January – June 2011). The changes in this case may have been to add redundancy to the C&C infrastructure as it grew.

These case studies provide three key insights into the life-cycle of mobile threats. First, in the case of threat ϵ , the malicious app was not publicly disclosed until months after its peak activity. In this case, reactive security measures failed to detect a threat until well after it was most effective. Second, both of these threats show a growth pattern similar to those shown in non-mobile malware studies [8]. Third,

⁴Only the top four ASes are presented; these comprise all hosts for threat ϵ and 308 out of the 316 hosts used by threat β .

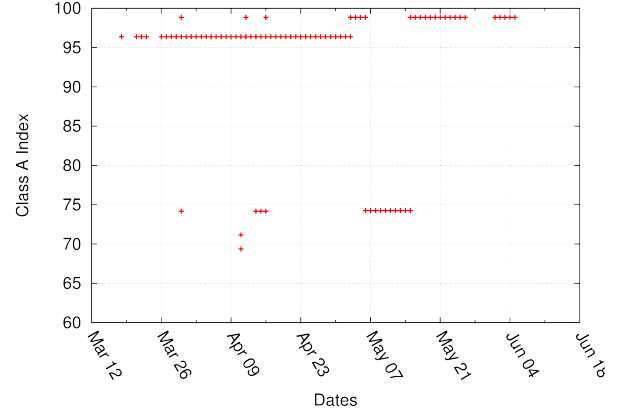


Figure 12: Threat ϵ 's host infrastructure shows agility comparable to non-mobile botnets

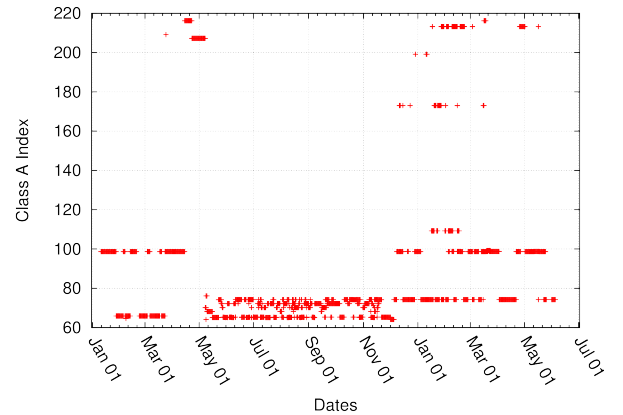


Figure 13: Threat β 's host infrastructure also shows agility comparable to non-mobile botnets

the agility in the hosting infrastructure used by these threats does not resemble professional DNS hosting. Rather, they are similar to non-mobile botnet operators that commonly use tactics like moving to hosts in different networks and countries to provide agility to their illicit operations. In future work, the agility seen in these mobile threats may be exploited by traditional DNS reputation systems to detect potentially suspicious domain names in the mobile space.

6 Related Work

The importance of mobile networks is increasing as society becomes more reliant on mobile devices such as smartphones, tablets, and mobile broadband cards. Several works have examined mobile device network traffic to learn about the general network characteristics of those

devices [19, 24, 18]. Past studies have shown that certain design considerations have made these networks inherently vulnerable to Denial of Service (DoS) attacks. Traynor et al. [41, 42] proposed a text messaging DoS with only the bandwidth of a cable modem. This research demonstrated a growing class of vulnerabilities due to the increasingly intertwined connectivity between the Internet and traditional voice networks. Other work has shown that the use of data communication protocols on voice networks creates the potential for failure under modest loads [35, 44, 43, 40]. Accordingly, significant effort has now been directed towards the analysis of potentially malicious mobile applications.

Numerous studies have highlighted the weaknesses and potential for misuse of various aspects of the Android security model [16, 23, 20, 22, 12]. Other work on Android devices suggests that it is difficult to tell if an application breaks any phone-wide security policies [17] and has resulted in tools to aid in the analysis of Android applications [14, 15]. Additional studies have surveyed the types of malware seen in the wild and evaluated the efficacy of different techniques in preventing and identifying such threats in the future [21]. However, app analysis alone provides an incomplete picture of the current state of malware on mobile devices and networks.

Network level analysis of malicious behavior offers a complementary means of characterizing and mitigating malware. For example, a popular method of preventing or limiting the spread of malware is the use of Internet blacklists. IP blacklists provide a list of known bad actors in the form of IP addresses which network operators can subsequently block; however, the use of DNS to build malicious network infrastructures has grown due to its resilience against IP blacklisting [36, 37]. Consequently, a significant amount of work has focused on analyzing those networks at the DNS level [27, 45, 29, 13, 25]. This has led to the creation of systems that are able to detect malicious domains through the use of passive DNS monitoring and machine learning [7, 10]. Furthermore, recent work has shown that detection of malicious domains can also be accomplished by passively monitoring DNS at the upper levels of the DNS hierarchy; this allows DNS operators to independently detect malicious domains without relying on local recursive DNS servers [8]. Ultimately, these systems allow network operators to assemble DNS blacklists of malicious and suspicious domains in order to detect and prevent malicious activity on the network.

Though there has been considerable effort targeted towards detecting network malware, it has been focused primarily on traditional wired networks. The question of whether such threats differ or even exist in real mobile networks has yet to be evaluated through empirical results.

7 Conclusions

In this paper, we presented a study of traffic obtained from a major US cellular provider as well as a major US non-cellular Internet service provider. Our work provides an in-depth understanding of the Internet infrastructure used for mobile malware. In particular, we showed that the network infrastructure used by mobile applications is part of the core Internet infrastructure used by applications in the non-cellular world; in other words, the mobile web is part of the Internet. We presented evidence showing that the mobile malware discovered by the research community appears in a minuscule number of devices in the network; this suggests that mobile application markets are already providing adequate security for a majority of mobile devices. We compared traffic to suspicious hosts between different mobile device platforms and demonstrated that iOS devices are no less likely than other platforms to reach out to such devices. Finally, we analyzed two major mobile threats and found that their network characteristics are similar to those of non-cellular botnets. Overall, these findings suggest that there are commonalities, in terms of both network infrastructure and characteristics, between malicious mobile applications and non-cellular malware. Therefore, we should leverage our successful experiences with DNS monitoring and reputation systems for non-cellular ISPs to develop a similar system for cellular carriers to identify (emerging) mobile threats. We leave this as a future work.

8 Acknowledgements

This work was supported in part by the US National Science Foundation (DGE-1148903, CNS-0916047, CAREER CNS-0952959, TWC-1222699, 0831300) and the Office of Naval Research (N000140710907, N000140911042). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the Office of Naval Research.

References

- [1] Admob. <http://www.admob.com>.
- [2] Heytell - instant voice messaging. <http://heyttell.com/front.html>.
- [3] Android/DroidKungFu.A!tr. http://www.fortiguard.com/encyclopedia/virus/android_droidkungfu.a!tr.html, June 2011.
- [4] Android/FakeDoc.A!tr. <http://www.fortiguard.com/av/VID3304615>, Dec. 2011.
- [5] Android/Steek.A!tr. <http://www.fortiguard.com/av/VID3458224>, Jan. 2012.
- [6] Alexa. The web information company. <http://www.alexa.com/>, 2007.

- [7] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster. Building a dynamic reputation system for DNS. In *Proceedings of the 19th USENIX Security Symposium*, 2010.
- [8] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou, and D. Dagon. Detecting malware domains at the upper DNS hierarchy. In *Proceedings of the 20th USENIX Security Symposium*, 2011.
- [9] BBC News. Malicious app penetrates iTunes store to test security. <http://www.bbc.co.uk/news/technology-15635408>, 2012.
- [10] L. Bilge, E. Kirder, C. Kruegel, and M. Balduzzi. Exposure: Finding malicious domains using passive DNS analysis. In *Proceedings of the 18th Network and Distributed Systems Symposium*, 2011.
- [11] Cathal Mullaney and Jeet Morparia. Android.Tonclank. http://www.symantec.com/security_response/writeup.jsp?docid=2011-061012-4545-99, June 2011.
- [12] E. Chin, A. Felt, K. Greenwood, and D. Wagner. Analyzing inter-application communication in Android. In *Proceedings of International Conference on Mobile Systems, Applications, and Services*, 2011.
- [13] D. Dagon, M. Antonakakis, P. Vixie, T. Jinmei, and W. Lee. Increased DNS forgery resistance through 0x20-bit encoding: security via leet queries. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*, 2008.
- [14] W. Enck, P. Gilbert, B. Chun, L. Cox, J. Jung, P. McDaniel, and A. Sheth. Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX Symposium on Operating Systems Design and Implementation*, 2010.
- [15] W. Enck, D. Octeau, P. McDaniel, and S. Chaudhuri. A study of Android application security. In *Proceedings of the 20th USENIX Security Symposium*, 2011.
- [16] W. Enck, M. Ongtang, and P. McDaniel. On Lightweight Mobile Phone Application Certification. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2009.
- [17] W. Enck, M. Ongtang, and P. McDaniel. Understanding Android security. *IEEE Security and Privacy Magazine*, 7(1):50–57, 2009.
- [18] J. Erman, A. Gerber, K. K. Ramadrihnan, S. Sen, and O. Spatscheck. Over the top video: the gorilla in cellular networks. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, IMC '11, page 127136, New York, NY, USA, 2011. ACM.
- [19] H. Falaki, D. LyMBERopoulos, R. Mahajan, S. Kandula, and D. Estrin. A first look at traffic on smartphones. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, IMC '10, page 281287, New York, NY, USA, 2010. ACM.
- [20] A. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *Proceedings of the ACM Conference on Computer and Communication Security*, 2011.
- [21] A. Felt, M. Finifter, E. Chin, S. Hanna, and D. Wagner. A survey of mobile malware in the wild. In *ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, 2011.
- [22] A. Felt, K. Greenwood, and D. Wagner. The effectiveness of application permissions. In *Proceedings of the USENIX Conference on Web Application Development*, 2011.
- [23] A. Felt, H. Wang, A. Moschuk, S. Hanna, and E. Chin. Permission re-delegation: Attacks and defenses. In *Proceedings of the 20th USENIX Security Symposium*, 2011.
- [24] A. Gember, A. Anand, and A. Akella. A comparative study of handheld and non-handheld traffic in campus Wi-Fi networks. In *Proceedings of the 12th international conference on Passive and active measurement*, PAM'11, page 173183, Berlin, Heidelberg, 2011. Springer-Verlag.
- [25] S. Hao, N. Feamster, and R. Pandrangi. An internet wide view into DNS lookup patterns. Technical report, Verisign Labs, 2010.
- [26] T. Holz, C. Gorecki, K. Rieck, and F. Freiling. Measuring and detecting fast-flux service networks. In *Proceedings of Annual Network and Distributed Systems Security Symposium (NDSS)*, 2008.
- [27] J. Jung, E. Sit, H. Balakrishnan, and R. Morris. DNS performance and the effectiveness of caching. *IEEE/ACM Transactions on Networking*, 10(5):589–603, 2002.
- [28] M. Kassner. Google Play: Android's Bouncer can be pwned. <http://www.techrepublic.com/blog/security/google-play-androids-bouncer-can-be-pwned/8053>, 2012.
- [29] C. Liu and P. Albitz. *DNS and BIND*. O'Reilly Media, 5th edition edition, 2006.
- [30] J. Ma, L. Saul, S. Savage, and G. Voelker. Beyond blacklists: Learning to detect malicious web sites from suspicious urls. In *Proceedings of the SIGKDD Conference*, 2009.
- [31] P. Mockapetris. Domain names - concepts and facilities. RFC 1034 (Standard), Nov. 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936.
- [32] P. Mockapetris. Domain names - implementation and specification. RFC 1035 (Standard), Nov. 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2845, 3425, 3658, 4033, 4034, 4035, 4343, 5936, 5966.
- [33] Nino Fred Gutierrez. Android.Gonesixty. http://www.symantec.com/security_response/writeup.jsp?docid=2011-093001-2649-99, Sept. 2011.
- [34] Piotr Krysiuk. Android.Gttracker. http://www.symantec.com/security_response/writeup.jsp?docid=2011-062208-5013-99, June 2011.
- [35] F. Ricciato. Unwanted traffic in 3g networks. *ACM SIGCOMM Computer Communication Review*, 36(2):53–56, 2006.
- [36] P. Royal. Analysis of the kraken botnet. Technical report, Damballa Labs, 2008.
- [37] S. Shevchenko. Srizbi's domain calculator, 2008.
- [38] Takashi Katsuki. Android.Walkinwat. http://www.symantec.com/security_response/writeup.jsp?docid=2011-033008-4831-99, Mar. 2011.

- [39] Tim Wyatt. DroidDreamLight, new malware from the developers of DroidDream. <http://blog.mylookout.com/blog/2011/05/30/security-alert-droiddreamlight-new-malware-from-the-developers-of-droiddream/>, May 2011.
- [40] P. Traynor, C. Amrutkar, V. Rao, T. Jaeger, P. McDaniel, and T. La Porta. From Mobile Phones to Responsible Devices. *Journal of Security and Communication Networks (SCN)*, 4(6):719 – 726, June 2011.
- [41] P. Traynor, W. Enck, P. McDaniel, and T. L. Porta. Exploiting open functionality in sms-capable cellular networks. *Journal of Computer Security*, 16(6):713–742, 2008.
- [42] P. Traynor, W. Enck, P. McDaniel, and T. L. Porta. Mitigating attacks on open functionality in sms-capable cellular networks. *IEEE/ACM Transactions on Networking*, 17(1):40–53, 2009.
- [43] P. Traynor, M. Lin, M. Ongtang, V. Rao, T. Jaeger, T. L. Porta, and P. McDaniel. On cellular botnets: Measuring the impact of malicious devices on a cellular network core. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2009.
- [44] P. Traynor, P. McDaniel, and T. L. Porta. On attack causality in internet-connected cellular networks. In *Proceedings of the 16th Network and Distributed Systems Symposium*, 2007.
- [45] D. Wessels, M. Fomenkov, N. Brownlee, and K. Claffy. Measurements and laboratory simulations of the upper DNS hierarchy. In *Proceedings of Passive and Active Measurement Workshop*, 2004.
- [46] Yi Li. Android.Notcompatible. http://www.symantec.com/security_response/writeup.jsp?docid=2012-050307-2712-99, May 2012.
- [47] Y. Zhou and X. Jiang. Dissecting Android Malware: Characterization and Evolution. In *Proceedings of the IEEE Symposium on Security and Privacy (OAKLAND)*, 2012.
- [48] Y. Zhou, Z. Wang, W. Zhou, and X. Jiang. Hey, You, Get off of My Market: Detecting Malicious Apps in Official and Alternative Android Markets. In *Proceedings of the 19th Network and Distributed System Security Symposium (NDSS)*, 2012.