

Please Pay Inside: Evaluating Bluetooth-based Detection of Gas Pump Skimmers

Nishant Bhaskar, Maxwell Bland, Kirill Levchenko[†], and Aaron Schulman
University of California, San Diego [†]*University of Illinois Urbana-Champaign*

Abstract

Gas pump skimming is one of the most pervasive forms of payment card attacks in the U.S. today. Gas pump skimmers are easy to install and difficult to detect: criminals can open gas pump enclosures and hide a skimmer in internal payment wiring. As a result, officials have resorted to detecting skimmers by performing laborious manual inspections of the wiring inside gas pumps. In addition, criminals can also avoid being caught using skimmers: many gas pump skimmers have Bluetooth connectivity, allowing criminals to collect payment data safely from inside their car.

In this work, we evaluate if the use of Bluetooth in skimmers also creates an opportunity for officials to detect them without opening gas pumps. We performed a large-scale study where we collected Bluetooth scans at 1,185 gas stations in six states. We detected a total of 64 Bluetooth-based skimmers across four U.S. states—all of which were recovered by law enforcement. We discovered that these skimmers were clearly distinguishable from legitimate devices in Bluetooth scans at gas stations. We also observed the nature of gas station skimming: skimmers can be installed for months without detection, and MAC addresses of skimmers may reveal the criminal entity installing or manufacturing them.

1 Introduction

Payment card skimming attacks at gas pumps have reached alarming levels. In 2018, law enforcement officials recovered 972 skimmers from gas pumps in Florida [11] and 148 skimmers from Arizona [10] alone. Based on industry estimates, a single skimmer can capture 30–100 credit cards per day [5] and each card, based on estimates from law enforcement officials, nets the criminal an estimated \$500 [6], resulting in a daily loss of \$15,000–50,000 per day of operation for each skimmer.¹ Less is known about how long a skimmer remains in operation, but allowing for even one day

of operation per skimmer, 2018 losses exceed \$16 million across these two states.

Gas pumps are an ideal skimming target. Gas pumps have relatively weak security: their payment circuitry can be accessed with universal keys or crowbars, and reading payment data is as easy as tapping into a ribbon cable (Section 2.1). Gas pump skimmers can be hidden inside of a gas pump enclosure, making them difficult to detect. As a result, inspectors have resorted to manually opening the pumps to inspect their wiring for skimmers. Gas pump skimming has become so pervasive that the Arizona Department of Agriculture, Weights and Measures Division (AZWMSD) now checks for skimmers while doing routine inspections.² From 2016 to 2018, the AZWMSD looked for skimmers in 7,325 gas station inspections. Inspectors found skimmers in only 1.5% of these inspections.

Unfortunately, Law Enforcement (LE) rarely catch criminals while they are collecting payment data from gas pump skimmers. The reason is, many gas pump skimmers are equipped with Bluetooth connectivity [26, 27, 28, 29]. This allows criminals to remain in their car while wirelessly retrieving card payment data. While Bluetooth is a vital tool for criminals to exfiltrate data from gas pumps, it also could be an opportunity to make it easier to detect skimmers.

In this paper, we evaluate the effectiveness of detecting skimmers with Bluetooth scanning from a smartphone. Indeed, if a skimmer can be detected with a smartphone, then authorities can discover and remove skimmers passively and quickly while they visit a gas station for other reasons. We built a smartphone application to perform this study, called Bluetana. Bluetana collects all Bluetooth scan data that is available via the Android Bluetooth APIs. We equipped 44 volunteers in six U.S. states with smartphones running Bluetana. Our volunteers have collected scans at 1,185 gas stations, where they observed a total of 2,214 Bluetooth devices. In these scans, Bluetana detected a total of 64 skimmers installed at gas stations in Arizona, California, Nevada,

¹In Section 2.2, we compare these quoted estimates to other sources, and find them to be in agreement.

²For example, the “Vapor Recovery Inspection Pre-Test Checklist” has a checkbox for “Checked for Skimmers”.

and Maryland, and it was the sole source of information that led law enforcement to find 33 skimmers.

The primary result of this study is the first comprehensive look at how skimmers can appear in Bluetooth scans. Namely, we observe that it is feasible to differentiate skimmers from other common Bluetooth devices that appear in Bluetooth scans at gas stations (e.g., vehicle telemetry collectors). The main differentiating factor for the skimmers we observed, is that the Bluetooth Class-of-Device—a parameter not collected by any consumer Bluetooth scanning applications that we are aware of—is “Uncategorized”. We also find that signal strength is a reliable way to determine if a Bluetooth device is located near a gas pump, and thus could be a skimmer.

Our study reveals several problems with consumer Bluetooth-based skimmer detection applications [46, 2, 51]: (1) there are many legitimate products that appear at gas stations that use the same Bluetooth modules as known skimmers; therefore, MAC address-prefix based detection may lead to false positives, (2) there are many Bluetooth modules used in skimmers that do not comply with IEEE MAC assignment requirements. We also debunk advice on how to find skimmers with Bluetooth scans from authorities [4] and viral information from social media [33]. For instance, none of the skimmers we found using Bluetooth scans have a name that is a long string of letters and numbers.

Performing this in-depth study brought to light several important operational lessons learned about the importance of detecting skimmers with Bluetooth. Using Bluetooth scans, officials detected skimmers while driving by gas stations that they otherwise would not have inspected. We also witnessed several instances where an inspector tried to find a skimmer, but could not find it on their first pass looking inside a gas pump. However they persisted and found it based on the knowledge that a suspected skimmer had appeared in Bluetooth scans. Surprisingly, we observed that there are skimmers installed in the same gas station, or city, that have very similar MAC addresses—indicating their source is a single criminal entity. We even found skimmers installed hundreds of miles away that had surprisingly close MAC addresses.

The rest of the paper is organized as follows: Section 2 provides background on internal gas pump skimming: their construction, monetary incentive, and prevalence in the wild. Section 3 is an overview of our large-scale Bluetooth scan collection methodology. In Section 4, we present the results of our study: what the skimmers we detected look like, how they compare to skimmers recovered independently by Law Enforcement, and whether they are well hidden in the Bluetooth environment. In Section 5, we present possible counter measures to the Bluetooth detection. In Section 6 we present the operational lessons we learned about skimming and criminal investigation procedure, while performing our large scale measurement study. Section 7 is related work, and we conclude in Section 8.



Figure 1: An internal Bluetooth-based skimmer wrapped in grey tubing to blend in with the cabling inside the fuel pump. This skimmer was detected by Bluetana in Tempe, AZ.

2 Background

Skimmers are illicit devices that capture credit card magnetic stripe data when a card is used at a point-of-sale (PoS) terminal or automatic teller machine (ATM). External skimmers use a magnetic head concealed in a false faceplate to read the magnetic stripe of a card as it is inserted into the real card reader. However, this paper is concerned with a newer class of skimmers, called *internal skimmers*, that are installed entirely inside a PoS terminal or ATM, leaving no visual evidence of its presence [47]. Internal skimmers are attached inline to the cable that connects the card reader to the main circuit board of the PoS terminal, tapping into the data and drawing power. To make data collection easier, many internal skimmers include a Bluetooth-to-serial module that allows the perpetrator to covertly collect the “skimmed” card data from a safe distance. These skimmers are built using commodity hardware with a total unit cost of \$20 or less.

Fuel pumps with a built-in PoS terminal have become a very popular target for such internal skimmers: they are unattended, easy to access, and have poor physical security, which make it easy to install a skimmer without being noticed. In a typical installation scenario, an attacker positions a van at a fuel station to block the station attendant’s view of the target pump (Excerpt in A.2), opens the fuel pump using a common master key or crowbar, and clips a discreet gumstick-sized skimmer to the ribbon cable between reader and main circuit board using a vampire clip (Figure 1). The entire process to install skimmer can take less than 10 seconds [1]. The perpetrator can then return to the station with a smartphone, and without leaving their vehicle, connect to the skimmer using Bluetooth and download the card data.

2.1 Internal Bluetooth Skimmers

The subject of our study are *internal, Bluetooth-based skimmers* that are installed in fuel pump PoS terminals. Figure 2 shows a typical Bluetooth skimmer, recovered from a fuel station in Southern California. This skimmer consists of a “Teensy” development board with an ARM Cortex-M4F micro-controller and a Roving Networks RN-42 Bluetooth-to-serial module. It also includes connectors for tapping into the wiring inside the pump (not shown).

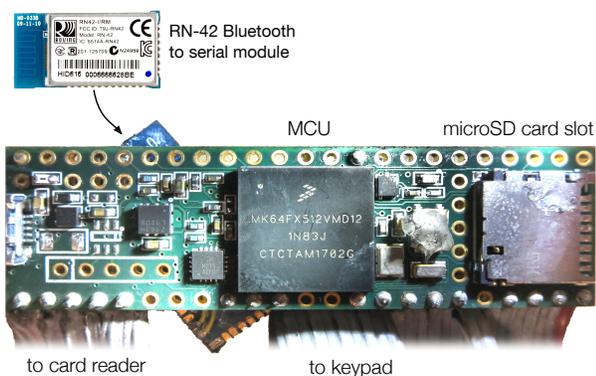


Figure 2: Parts of a typical internal Bluetooth-based fuel pump skimmer. This skimmer was detected by Bluetana.

Connections. In the figure, the ribbon cable on the left intercepts or replaces the ribbon cable that connects the magnetic stripe reader to the PoS terminal main board. The skimmer also uses this connection for power: the power and ground pins of the Teensy (on far left of board, not visible in Figure 2) are connected to power and ground on the card reader cable. The ribbon cable on the right intercepts or replaces the ribbon cable from the PoS keypad. This allows the perpetrator to capture additional card verification data, namely the debit card PIN or credit card billing ZIP Code. Availability of a PIN code with a stolen debit card in particular, can increase its value five-fold on the black market (Table 1). However, not all skimmers capture keypad data.

Most gas station skimmers read the unencrypted data pulled from magnetic stripe readers. Card issuers feel that removing sensitive data from the magnetic stripe on cards will help to solve the problem [42]. Newer literature has demonstrated attacks on chip payment systems [13, 15], and law enforcement in Latin America have begun to find EMV skimmers that are Bluetooth enabled [30, 3].

Controller board. The skimmer pictured in Figure 2 used a Teensy micro-controller development board equipped with a 120 MHz ARM Cortex-M4F micro-controller made by Freescale Semiconductor. By using a development board, a skimmer requires only rudimentary electronic assembly: soldering wires to the development board.

However, skimmers have also been found using what appeared to be fully custom-designed boards. These are compact, making them better for hiding in the dispenser. Examples of micro-controllers used in recovered skimmers include Microchip PIC18F4550 [2] and Atmel XMEGA128A4U [3].

Storage. The Teensy board also has a microSD card slot for additional data storage. Skimmers built on custom PCBs have also used flash and EEPROM ICs for storage. The storage capacities vary across designs, with examples using the PCT25VF032B (32-Mbit) [3] and M25P16VP (16-Mbit) [2].

Bluetooth module. The skimmer shown in Figure 2 uses a Roving Networks RN-42 module, an inexpensive Bluetooth-to-serial module found in many skimmers. In Section 3.1 we describe characteristics of popular Bluetooth-to-serial modules used in recovered skimmers for wireless data exfiltration. On the Bluetooth side, a Bluetooth-to-serial module provides a Serial Port Peripheral interface, which most operating systems recognize as a Bluetooth modem and instantiate a serial device for it. Operating systems will create a corresponding serial device, allowing user-space applications, namely a criminal’s card dumping application, to communicate with the module. On the hardware side, a Bluetooth-to-serial module provides a TTL-level receive and transmit pin, allowing it to interface to any micro-controller UART. The module this allows even the simplest micro-controller to communicate via Bluetooth with a host device. The 2.4GHz Bluetooth antenna is included on the module’s circuit board (exposed area to the left of the metal shield for the module shown in Figure 2), so the antenna is also hidden.

Bluetooth-to-serial modules generally require no configuration, however, most can be reconfigured using Hayes-style modem AT commands. In Section 4.1 we describe the configuration capabilities of popular modules. Notably, all of the Bluetooth-to-serial modules we found in skimmers support changing the device MAC address, Bluetooth device name, changing the pairing password, and the ability to become non-discoverable once paired.

2.2 Economics of Carding

Stealing and monetizing stolen credit and debit card data, called *carding* by its practitioners, is a well-studied form of financial fraud, however, reliable estimates of losses resulting from a single skimmer are difficult to find. To the criminal operating a skimmer, the expected revenue per skimmer breaks down as:

$$W = (\text{card value}) \times (\text{cards per day}) \times (\text{days deployed}).$$

Of these, we found published estimates for only the first two quantities, and very little about skimmer lifetimes. Here, we summarize the available data with the goal of estimating the losses incurred by a single skimmer.

Card value. To monetize stolen credit card data, skimmer installers have two options: sell the data on the black market, or cash out the cards on themselves. Based on our survey of sites selling stolen card data, black market prices for stolen cards fall in the \$10–220 range, depending on whether the card is a debit or credit card, and whether it comes with a PIN (for debit) or billing ZIP code (for credit). Table 1 provides a summary of these prices with references.

Criminals can also cash out the cards themselves. Debit cards with a PIN are often cashed out by withdrawing money from an ATM, while credit cards are often cashed out by

<i>Scheme</i>	<i>Value</i>	<i>Reference</i>
Black market price		
Debit, no PIN	\$20–30	[35, 49, 21, 44]
Debit with PIN	\$110–220	[31, 49, 44]
Credit, no ZIP	\$10–25	[35, 49, 21, 44]
Credit with ZIP	\$25–60	[35, 49, 21, 44]
Cash-out value		
Credit or Debit (standard)	\$400–800	[19, 40, 18, 56]
Credit (premium)	\$1,000	[40, 45, 20]
Bank and merchant loss		
Credit	\$1,003	[1]
Debit	\$650	[12]
Consumer liability		
Debit (> 60 days)	unlimited	15 USC 1693g
Debit (< 60 days)	max \$500	15 USC 1693g
Debit (< 2 days)	max \$50	15 USC 1693g
Credit	max \$50	15 USC 1643
Prosecuted loss		
Credit or debit	\$500	[6]
Court documents		
Credit	\$362–400	[36, 16, 8, 7]
Debit	\$665–1132	[9, 50]

Table 1: Value of stolen credit and debit cards.

purchasing high-value merchandise (e.g. iPhones) and reselling them. Reported cash-out values for debit and credit cards range between \$400 and \$1,000, depending on credit limit associated with the card. We also conducted a survey of cash-out values reported in court documents involving skimmers.³ Several cases reported specific cash-out values, rather than ranges. The debit card cash-out values were \$1132 [36], \$444 [16] \$665 [8], \$1354 [7]. The credit card cash-out values were \$362 [50] and \$400 [9].

Losses due to credit and debit card fraud are borne largely by banks and merchants. This is likely because consumer liability for fraud in the U.S. is limited to \$50 for credit cards, and \$50 or more for debit cards (depending on how quickly the consumer reports the fraud). Industry estimates for losses per-card incurred by banks are \$650 for debit cards and, \$1,003 for credit cards [1, 12]. The U.S. Sentencing Commission estimates per-card losses at \$500 or more.

Cards per day. The number of cards a skimmer captures each day depends on the number of transactions at that pump, which will vary by station. Ripplshot, a payment fraud prevention service, states: “a single compromised pump can capture data from roughly 30–100 cards per day” [5]. The lower end Ripplshot’s estimate agrees with the estimate of 20–50 cards per day we received from U.S. law enforcement agents. In addition, we found two court documents that report criminals captured 25 [9] and 30 [8] cards per day. We

³We surveyed only documents available without fee from Court Listener.

<i>Location & Year</i>	<i>Recovered skimmers</i>	<i>Skimmed stations</i>	<i>Skimmers / station</i>	<i>Skimmers / 10⁶ people</i>
San Diego				
FY 2018	42	11	3.2	11.9
Arizona				
2016	88	54	1.6	4.3
2017	57	46	1.2	2.7
2018	148	86	1.7	6.9
All	293	134	2.2	14.0
Florida				
2016	207	162	1.3	10.0
2017	650	432	1.5	31.1
2018	972	524	1.8	45.6
All	1,829	1,029	1.7	87.4

Table 2: Prevalence of skimming in three regions of the U.S.

also studied 10 skimmers recovered from the field, which we were told were used and wiped daily. We found an average of 20 cards per skimmer, divided evenly between debit and credit cards.⁴

Days deployed. Internal skimmers are not limited by battery life and can remain in operational indefinitely, because they draw power from the PoS circuitry, Skimmer lifetime, then, is limited only by how long they can remain undetected. Unfortunately, there is little reliable data on this. Our only direct experience is our discovery of a pair of skimmers that remained undetected for six months (Section 3.1). However, LE informed us that criminals may leave skimmers in gas pumps after only a few days of retrieving card data and moving on to another location. Given the very limited data available on skimmer lifetimes, we instead consider skimmer value *per day of operation*.

Cashout success rate. Our analysis of court documents revealed that criminals are often unsuccessful when trying to cashout a skimmed card. This may be due to a variety of reasons, such as the following: incorrectly reading card data, hitting daily withdrawal limits, and activating fraud alerts. Several cases mentioned that criminals were not successful in cashing all skimmed cards. One case mentions a specific cashout success rate of 47% [7].

Total skimmer value. Finally, we estimate the range of per-day revenue from a skimmer based on the prior figures. Our low end estimate is \$4,253 (25 cards per day, cashout of \$362 per card, and 47% cashout success rate), and our high end estimate is \$63,638 (100 cards per day per day, \$1,354 cashout per card, and cashout success rate of 47%).

⁴These skimmers were provided to us because they were removed by the station owner, rather than LE, making them unsuitable for use as evidence.

2.3 Skimmers Recovered in the Wild

To understand the prevalence of skimmers in the wild, we obtained data on recovered skimmers from three regions in the United States: San Diego and Imperial counties of California, with a combined population of 3.5 million; the state of Arizona, with a population of 7 million inhabitants; and the state of Florida, with a population of 21 million inhabitants. Table 2 summarizes the statistics. We note that these numbers do not represent *all* recovered skimmers. For San Diego and Imperial counties, our statistics represent the number of skimmers found by or reported to a U.S. federal law enforcement agency. For Arizona and Florida, our statistics represent skimmers found by or reported to the AZWMSD and the Florida Department of Agriculture and Consumer Services.

The number of recovered skimmers has increased from 2016 to 2018 in both Florida and Arizona. The total number of skimmers recovered in 2018 across the three geographic regions is significant: if each skimmer operated for just one day, we estimate their total monetary impact would be \$17.43 million. Yet, as the skimmers-per-million people number shows, the possibility of an average consumer encountering a skimmer at a gas station is quite small.

3 Data Collection Methodology

Driven by the observation that skimmers are hard to find—few pumps in San Diego, Arizona, and Florida have been found to have skimmers installed in them (Table 2)—we created a tool, called Bluetana, to evaluate the effectiveness of Bluetooth-based skimmer detection. We begin by presenting an overview of the tool and the data it collects. Then we describe how Bluetana identifies suspicious devices and directs users to collect additional data. Finally, we discuss how we retroactively inspect data to find skimmers.

3.1 Crowdsourcing Bluetooth Scanning

We developed Bluetana, an Android-based measurement tool that officials and volunteers use to scan for skimmers at gas stations. Bluetana scans for nearby Bluetooth—both Classic and Bluetooth Low Energy (BLE)—devices every 5 seconds using Android’s Bluetooth API. It collects the Bluetooth scans and geo-location data, and uploads this data to a secure database over a cellular link. Bluetana collects all of the Bluetooth scan data that Android makes available, including Device name, MAC Address, Class-of-Device⁵, and signal strength (RSSI).

How we visited 1200 gas stations. We outfitted 44 volunteers and inspectors in six U.S. states (CA, AZ, MD, NC, NV, IL) with low-end smartphones running Bluetana in kiosk

⁵Class-of-Device is twenty four bits indicating the device’s intended use, such as *smartphone* or *speaker*.

mode (they could not close the application). We selected officials who frequent gas stations as part of their daily job duties. Primarily, they were Weights and Measures inspectors.

Indicating suspicious devices to inspire data collection

The Bluetana display shows a list of Bluetooth devices detected during scanning. When Bluetana detects a potential skimmer, it indicates this to the user by highlighting the device record (Figure 4). The Bluetooth scan profile of the modules that have been found in skimmers inform which devices we highlight in Bluetana.

Skimmers recovered by LE are often found to use CSR (Qualcomm) chip-set-based Bluetooth modules. Our highlighting procedure primarily looks for the default Bluetooth profile of these modules—with the exception of the Device Name which can be missing due to poor signal strength, and modified by criminals in an attempt to hide the device (Section 4). The factory default Bluetooth scan profile (i.e., MAC prefix, Device Name, and Class-of-Device) of these modules are as follows:

Mod.	MAC Prefix	Dev. Name	Class of Dev.
RN	00:06:66	“RBNT-*”	Uncategorized
HC	Various	“HC-05/06”	Uncategorized

Bluetana chooses a highlight color via a three-step decision process, depicted in Figure 3. First, the app checks the device’s class. All skimmers studied within this work, whether discovered by Bluetana or not, had a device class of *Uncategorized*. If the device class is not uncategorized, the data is saved for later analysis. The device’s MAC prefix is then compared against a “hitlist” of prefixes used in skimming devices recovered by law enforcement. If the device has a MAC that is not on this hitlist, it is unlikely to be a skimmer, and the app highlights the record yellow. Next, if the device name matches a common product using the same MAC prefix, the record highlights in orange. If all three fields (MAC prefix, Class-of-Device, and Device Name) indicate the device is likely to be a skimmer, Bluetana highlights the record in red. The highlighting procedure is the result of a year of refinements based on our experience finding skimmers in the field, and Bluetana includes a remote update procedure to account for these incremental changes.

This simple highlighting proved to be vital to our data collection. Red serves as a cue to perform signal strength localization: it directed our users to collect more samples of signal strength to determine if a device is located in the gas pump area—and is therefore likely to be a skimmer. In several cases, Bluetana highlighting a device in red was the only reason officials performed a manual skimmer inspections: out of the 64 skimmers we found, 33 were recovered because an official started an inspection only after noticing a device was highlighted in red in Bluetana.

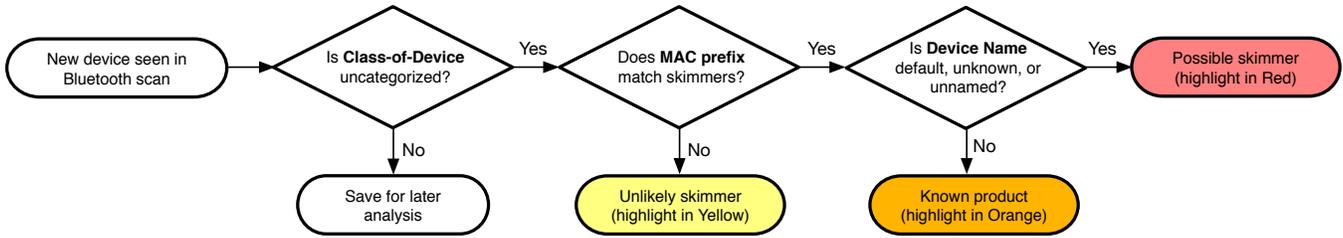


Figure 3: The procedure Bluetana uses for highlighting suspicious devices.

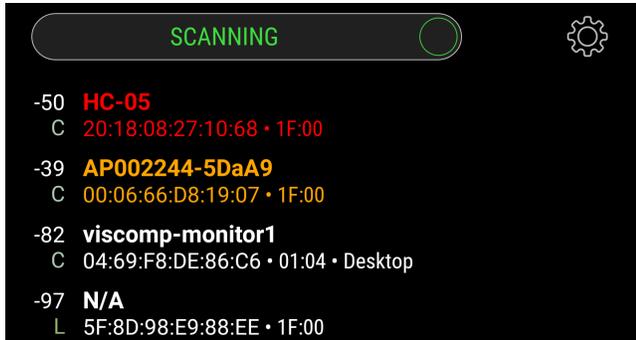


Figure 4: The Bluetana user interface. Bluetana highlights suspicious devices, inspiring users to collect more signal strength samples, and even perform inspections.

In one instance, an Arizona Weights and Measures inspector was driving by a gas station when two red highlighted devices appeared in Bluetana. He made an unscheduled stop at the gas station, performed a skimmer inspection, and discovered two skimmers. Figure 5 shows a portion of the official Arizona inspection report documenting this incident.

Bluetana’s highlighting procedure is more comprehensive than other skimmer detection apps on the Play Store. Scaife et al. [46] investigated the behavior of these apps and found that they flag skimmers based on either MAC prefix or Device Name. These apps would miss skimmers with non-standard MAC prefixes or customized (missing) device names which Bluetana was able to find (Section 4.1). Bluetana also found legitimate devices that would be considered skimmers by these apps (Section 4.2).

Identifying skimmers after data collection

During the study, we manually examined every Classic Bluetooth device observed at a gas station visit in real time (as Bluetana users upload their scan data). At the beginning of our study, we relied primarily on the signal strength of the device to determine if it was a suspected skimmer. By the nature of being installed inside a gas pump, the Bluetooth signal of a skimmer is strongest in the pump area. Other devices that we suspected to be skimmers all had a low signal strength in the pump area, because aside from the cars

parked at the pumps, the only places where a Bluetooth device would be located in the pump area would be inside the pump. Combining the signal strength and geo-location with satellite imagery of the gas station, we were able to easily detect when the signal was emanating from inside of a gas pump (example shown in Figure 6). While at a gas station, Bluetana users also noticed this by moving toward the pump area to see if the device’s signal strength increases.

If we saw any suspicious devices in the dataset, we alerted officials that they should inspect the pumps at the station in question. Initially, we did not know which of these devices were skimmers: many initial inspections we requested turned up empty handed. However, as the study progressed, we improved our understanding of the profile of skimmers.

A natural experiment observing deployment duration

Having a database of all prior scans made it possible for us to look for skimmers that we may have missed in the past. In particular, looking back in at the database led to us to discover two skimmers that we had initially missed. A retroactive analysis of two stations discovered skimmers that were still operating even though we first detected them *six months* earlier. This natural experiment is likely the first concrete data on how long skimmers can be installed without being found in a routine or complaint-induced pump inspection.

3.2 Limitations

Selection bias

We designed our data collection to look for a specific type of gas pump skimmer: one that uses a Classic Bluetooth module, and is discoverable in Bluetooth scans. Our contacts in LE confirmed that this type of skimmer has been found in gas stations across the entire U.S. They also reported that these skimmers are particularly common in Arizona and California; therefore, these states were the focus of our study.

The results of our study may not be representative of the nature of gas pump skimming across the country. Criminals in other regions may evade Bluetooth-based detection by using alternate exfiltration methods (e.g., Bluetooth Low Energy and SMS), or configurations (e.g., non-discoverable mode). We outline these countermeasures in Section 5.



BMF # [REDACTED]

INSPECTION # [REDACTED]

TEST DATE [REDACTED]

PAGE 1 OF 1

COMMENTS / NOTES

While using the "Bluetana" scanner two items showed up in red. I opened a fueling device skimmer inspection then announced myself to location staff. The scanner showed the strongest signal to the dispensers closet to [REDACTED] In dispensers 1/2 and 5/6 I found skimmers installed. For a total

Figure 5: Bluetooth scanning helps inspectors find more skimmers because they detect skimmers when driving by a gas station.

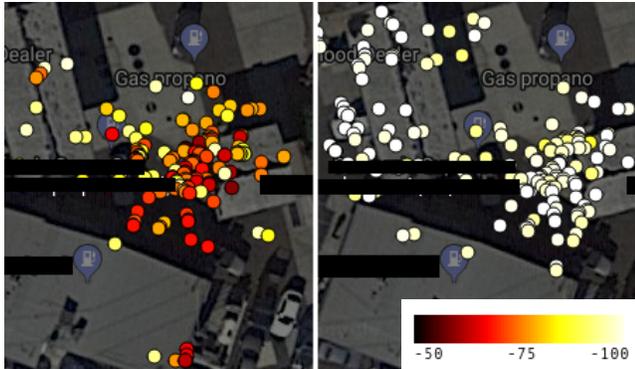


Figure 6: Combining RSSI data with satellite imagery reveals if a device is located in the pump area of a gas station.

Bluetana does not connect to devices

We could collect more data about Bluetooth devices by trying to connect to them. This could be useful for conclusively detecting a skimmer or collecting information about the type of Bluetooth device. By sending commands that skimmers are known to respond to, Bluetana would be able to see if the device responds equivalently to known skimmers. This is precisely what one of the current Bluetooth skimmer scanning applications on the Play Store does.

This practice may seem innocuous, but our discussions with law enforcement indicate that this could overwrite information critical to future investigations. The problem is, internal registers in many skimmer Bluetooth modules records the last-paired MAC address. This information can be used to link a suspect possessing a smartphone or laptop with their skimmers. The typical forensic evidence collection performed by law enforcement on skimmers includes collecting the last-paired MAC address [48].

4 Results

In this section, we present the results of our 19 month study of Bluetooth devices observed with Bluetana at 1,185 gas stations across six U.S. states (CA, AZ, NV, MD, IL, NC). During the course of this study, Bluetana detected 64 skimmers operating in 34 gas stations; all of the skimmers were

removed from the pumps by local and federal law enforcement agents. Bluetooth scanning is a surprisingly effective way of detecting skimmers: in Arizona, Bluetana has detected skimmers at 1.58% of the 491 stations it scanned, and routine inspections by state inspectors had a similar detection rate of 1.5% from 2016 to 2018.

The primary result of this study is as follows: there are distinct characteristics of the 64 internal skimmers detected by Bluetana that differentiate them from the 2,562 other Bluetooth devices that Bluetana found at gas stations (e.g., car stereos). Namely, these skimmers were predominately using the default Bluetooth module configuration. Additionally, we discovered that some criminals use a custom Device Name in an apparent attempt to hide their skimmers from Bluetooth scans. These custom Device Names stand out, making them easier to differentiate from other devices.

4.1 What Do Skimmers Look Like in Scans?

We begin by presenting how skimmers we observed appear in Bluetooth scans. We describe the properties of two sets of skimmers: 64 skimmers that we detected in the field during the course of this study, as well as 23 skimmers that were independently recovered by two LE agencies. The 23 skimmers recovered independently by LE have similar characteristics to the 64 that Bluetana detected in the field. The Bluetooth characteristics of these skimmers are detailed in Table 3. We now analyze the following properties: Class-of-Device, MAC prefix, and Device Name.

All of the skimmers are “Uncategorized” Class-of-Device

Class-of-Device is primarily used to select the icon that indicates the category of a device in a Bluetooth scan (e.g., Headphones). Bluetooth modules used in skimmers analyzed in this study (i.e., HC and RN), have an “Uncategorized” Class-of-Device assigned by default. Changing Class-of-Device on these modules is trivial: the modules provide a serial command to set it. Despite this, criminals do not appear to be modifying the Class-of-Device on any of the skimmers we observed: all of the 87 skimmers detected by Bluetana and recovered independently by LE used the default “Uncategorized” device class.

<i>Bluetooth Scan Property</i>	<i># of skimmers</i>	
	<i>Bluetana</i>	<i>LE</i>
Class-of-Device		
Uncategorized	64	23
Manufacturer (MAC prefix)		
Roving Networks		
00:06:66	45	13
Shenzhen Bolutek		
98:D3:31	1	
<i>Unknown</i>		
20:13:04	1	
20:17:11	1	
20:18:01	2	
20:18:04	1	
20:18:07	1	
20:18:08	4	10
20:18:09	4	
20:18:10	1	
20:18:11	2	
98:D3:35	1	
Device Name		
<i>Default</i>	36	23
[Law enforcement]	2	
[Mobile phone]	4	
[Indescript object]	2	
[Numerical]	2	
<i>Unnamed</i>	18	
Total	64	23

Table 3: Bluetooth scan properties of skimmers observed during our study. The exact Device Names are not shown, instead we describe the names we found.

MAC prefixes are often manufacturer defaults

Bluetooth module manufacturers burn a MAC address into the module’s EEPROM. Although it is possible to change the MAC with a SPI-based reprogramming of the CSR chip’s EEPROM, we have not observed any skimmers that have a modified MAC. The first three bytes (prefix) of the MAC address typically correspond to the manufacturer of the device.

Although MAC address prefixes are often assigned by IEEE (e.g., all of the RN Bluetooth modules have the same manufacturer MAC prefix) the HC modules have a wide variety of MAC prefixes. Of the HC modules we observed, only one has a MAC prefix assigned by the IEEE. This could make it significantly more difficult to detect an HC-equipped skimmer. However, looking at of the MAC prefixes of the skimmers that we observed, a clear pattern emerges: manufacturers appear to be burning module manufacture date into the first four bytes of the MAC address in the following format: YY:YY:MM:(DD).

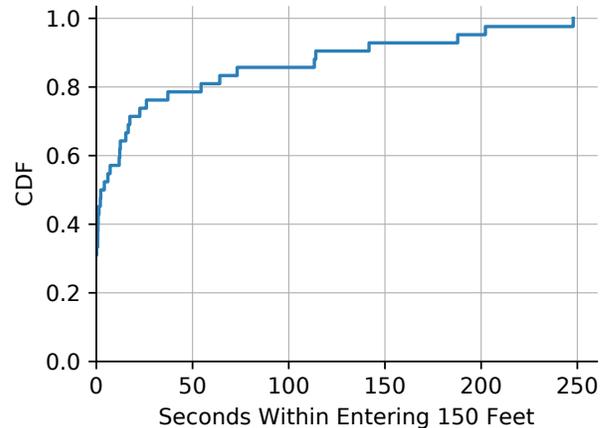


Figure 7: Skimmers are detected within a minute of passing near a gas station.

Device names are often default, occasionally customized

Device Names allow users to identify their devices in Bluetooth scans. They are assigned a factory default value by the manufacturers, and are modifiable by users. Most of the skimmers we observed had a default Device Name: namely, all of the skimmers provided by LE, and more than half the skimmers we detected in with Bluetana. A skimmer with a default Device Name looks innocuous, because some legitimate products using the same modules are also shipped with the default module name (Section 4.3). Occasionally, we found that criminals set a custom device name on their skimmers. This appears to be an attempt to make the skimmer look less suspicious. Bluetana detected custom-named skimmers with a variety of names. The custom names of skimmers discovered by Bluetana had variety: some were random strings of numbers, and others masqueraded as LE.

Bluetana did not detect a Device Name for several skimmers. This is expected because the device sends its MAC and Class-of-Device in the first scan response packet; it sends the device name in a subsequent packet (that may be missed).

Skimmers are detected within one minute

Bluetooth scanning has the benefit of detecting some skimmers without manually inspecting each of the pumps. However, attenuation from a gas pump’s metal enclosure, may limit the range that Bluetooth scans are effective. We analyzed the scans from Bluetana to see how long an official had to spend at a gas station before they detected the skimmers installed there (Figure 7). The median time to detection was 3 seconds, and 80% of the skimmers were detected within one minute. This is a 99% decrease in search time compared to the average of 30 minutes that inspectors take

State	Stations	Devices Observed			Days	Skimmers
		#	Avg.	Std.		
CA	571	1148	2.01	1.94	152	22
AZ	491	1140	2.32	2.03	130	36
NV	38	93	2.45	3.44	21	4
MD	23	42	1.83	1.86	14	2
IL	18	37	2.06	2.01	13	0
NC	10	20	2	1.67	10	0

Table 4: On average there are two Classic Bluetooth devices seen at each gas station; infrequently, there are skimmers.

to check a gas station for skimmers.⁶ This result indicates that inspectors can quickly stop at gas stations to check for Bluetooth-detectable internal skimmers.

4.2 Are Skimmers Distinguishable in Scans?

Next, we evaluate if the skimmers detected by Bluetana were clearly distinguishable from the other devices observed at gas stations. The primary result of this study is that these skimmers were not hidden well. Many of these skimmers use the default configuration of their Bluetooth modules. Legitimate devices using the same Bluetooth modules may have some default parameters, and a few have all of parameters set to the default. We conclude that by combining multiple characteristics: MAC prefix, Class-of-Device, and Device Name, there are only a small number of devices that could be confused with skimmers.

This study also reveals that when criminals creatively modify their skimmer’s Device Name, it makes detection easier. We also found that criminals could improve how they hide skimmers in Bluetooth scans. For example, they could change the Class-of-Device to hide as a more popular device (e.g., a smartphone).

Dataset Overview

Over the course of the 19 month study, Bluetana users visited 1,185 gas stations across six states (Table 4). During these visits, Bluetana detected a total of 64 skimmers—all of which were recovered by officials. These skimmers were in the presence of 2,214 other devices. On average, Bluetana saw 2.2 devices per station ($\sigma = 2.05$). Given that there are only a small number of Bluetooth devices seen per station, it may seem likely that these devices are all skimmers. However, only a small fraction (3.98%) of these devices matched the characteristics of the skimmers we observed during the course of our study.

We performed this study on Classic Bluetooth devices only. We did not include BLE because we are not aware of

⁶Source: discussions with inspectors.

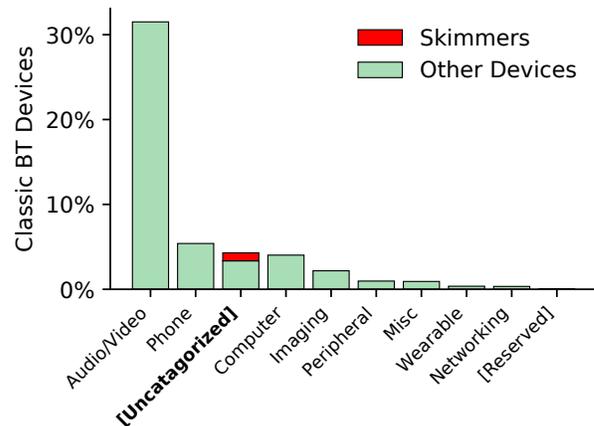


Figure 8: Skimmers appear in the third most common class of Bluetooth devices.

any internal gas station skimmers using BLE modules. However, we observed a large number of BLE devices at gas stations; therefore, switching skimmers to BLE modules may make them more difficult to detect with scanning tools like Bluetana (Section 5.1).

For this analysis, we only include the scan data that is collected the first time a Bluetana user visits a station. Restricting the dataset in this way ensures fairness in our results. Analyzing all inspections may bias our observation of what Bluetooth devices tend to be found at gas stations to those that were visited multiple times. Specifically, we only analyze scans performed the first time Bluetana is near a gas station (within 150 feet) for at least 30 seconds and up to 5 minutes. 22 out of 64 of the skimmers were detected on subsequent visits to gas stations, so they are not included in this analysis.

Skimmers are Uncategorized, but so are other devices

The only Bluetooth property that is common among all skimmers we observed is that they have an Uncategorized Class-of-Device. Figure 8 shows the distribution of Bluetooth device classes at gas stations. Uncategorized devices are the third most common Class-of-Device found by Bluetana (8.4% of devices). Although, out of the 1,185 gas stations that Bluetana users visited, Uncategorized devices were only observed at 143 gas stations (12.1%).

Other devices use the same modules as skimmers

Within the set of Uncategorized devices, we next look at the distribution of their MAC prefixes (Figure 9). We find that the Bluetooth modules used in skimmers are also used in many other legitimate devices. Specifically, more than half of the RN modules seen at gas stations were in skimmers, but

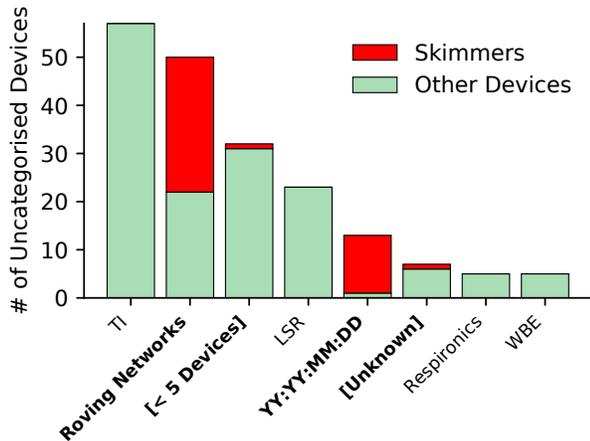


Figure 9: Many other devices appear to be using the same Bluetooth modules as skimmers.

there were many other devices that had RN modules. This is an important observation because a popular detection application, SkimPlus [51], only flags skimmers based on a hitlist of MAC prefixes [46]; it may incorrectly flag legitimate devices as skimmers.

The devices observed with MAC prefixes that were in the YY:YY:MM:DD format (likely HC modules) were mostly skimmers. There were many devices that had IEEE assigned MAC prefixes that were infrequently seen at gas stations (< 5 Devices). Only one of these devices was a skimmer. Also, there were many devices with MAC prefixes unknown to the IEEE, but not in the date format, only one of these devices was a skimmer. Overall, 102 devices out of 187 Uncategorized devices matched the MAC prefixes of Bluetana-observed skimmers. This reduces the number of stations where Bluetana detected skimmers to 79 out of the 143 stations where it found Uncategorized devices.

Default- and custom-named modules are often skimmers

Finally, we investigate if skimmers can be differentiated from other devices by their Device Name. The remaining 102 devices are Uncategorized and their MAC prefixes are either: Roving Networks, YY:YY:MM:DD, Unknown, or seen on less than five devices. Only 42 of these devices were confirmed to be skimmers.⁷ In Figure 10, we divide the remaining devices by their category of Device Name, including: *unnamed*, manufacturer *default*, known legitimate *product*, and *customized*. Devices observed by Bluetana with default names were often skimmers. Custom named devices were not common at gas stations but had a higher probability of being skimmers. Three skimmers were disguised as prod-

⁷We do not include 22 of the Bluetana-detected skimmers in this analysis because they were not detected on the first visit to a gas station.

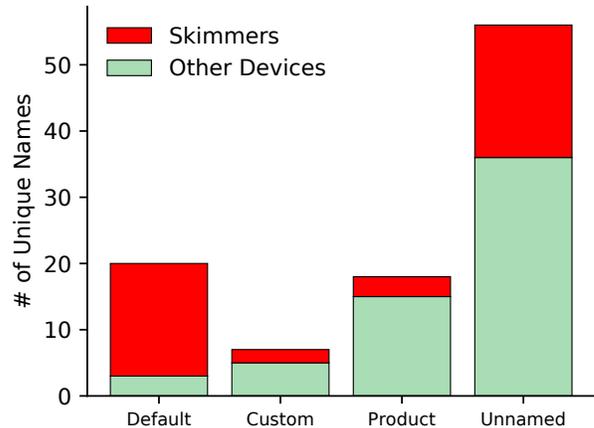


Figure 10: Default and custom names distinguish skimmers from legitimate devices.

ucts, however all three were distinguishable because their names were popular smartphones, which should not have the MAC prefix of Bluetooth-to-Serial modules. Bluetana missed capturing the Device Name for many of the skimmers, as well as other devices that it detected.

4.3 Accuracy of Bluetooth-based Detection

To evaluate the accuracy of Bluetooth-based detection, we analyze Bluetana scan data collected during inspections in Arizona. Specifically, there was a 7-month time period in which Bluetana was used by many of the Arizona inspectors (October 7, 2018 – May 7, 2019), and we compare the reports filed during these inspections with the scan data that Bluetana collected.

Missed skimmers

During this time period, there were 27 inspections where skimmers were found while an inspector was running Bluetana. A total of 42 skimmers were recovered during these inspections, of which Bluetana was able to detect 36. Therefore, Bluetana missed detecting 14.3% of the total skimmers recovered during these inspections.

We do not know exactly why Bluetooth-based scanning missed these skimmers. Half of the missed skimmers were from inspections where Bluetana detected other skimmers at the gas station. It is likely that these missed skimmers were not powered on due to improper installation. The remaining missing skimmers may have been built with alternate exfiltration methods, such as SMS [46], or even require physical recovery [47].

Incorrectly detected skimmers

Bluetana highlighted a device in red during 45 Arizona inspections where no skimmer was found. There were 757 total inspections where inspectors used Bluetana⁸, Bluetana may have incorrectly detect skimmers in 5.9% of inspections.

Incorrectly identifying skimmers is likely due to the fact that RN and HC modules are used in a variety of legitimate products, some of which are seen in and around gas stations. We found RN and HC modules in radar-based speed limit signs, weather sensors [38] automotive diagnostic scanners, scales [37] and fleet tracking systems [52]. Some of these devices have Device Names that clearly indicate what product they are, but would be confused with skimmers if the Device Name is missing. Unfortunately, several of these products also use the default Device Names on their Bluetooth modules (*RNBT-xxxx* or *HC-05*). These legitimate devices will look exactly like skimmers. In such cases, inspectors will need to rely on RSSI localization to determine if these devices are located inside a gas pump.

5 Countermeasures and Responses

This work is a single snapshot in an evolving landscape of attacks on payment systems. While Bluetana has proven effective at finding Bluetooth skimmers, it by no means represents the last move in the cat-and-mouse game. In the remainder of this section, we discuss what the next few steps in this arms race might look like. That is, given that inspectors and volunteers are using Bluetana, what can be the skimmer installers' next move, its cost, and what our response might be. It is possible for a determined and resourceful criminal to implement the countermeasures that we will be describing (particularly non-discoverable mode).

5.1 Switching to Bluetooth Low Energy

We have observed that by switching to BLE, criminals have *many* more places to hide. Figure 11 shows the cumulative distribution of the number of BLE and Bluetooth devices we saw at each fuel station. Under the filtering of Section 4, over 8,000 unique BLE devices were seen, making the ratio of Classic to BLE approximately 1:4.

Cost to attacker. There is almost no cost to criminals in switching their Bluetooth modules to BLE. In fact, newer EMV skimmers discovered in other countries are BLE enabled [30]. However, none of our contacts in law enforcement have encountered BLE-based gas station skimmers. It is possible that there is simply no incentive to switch: the same reason criminals have not yet adapted to masking their Bluetooth device class.

⁸This includes both routine and complaint/prior knowledge triggered inspections

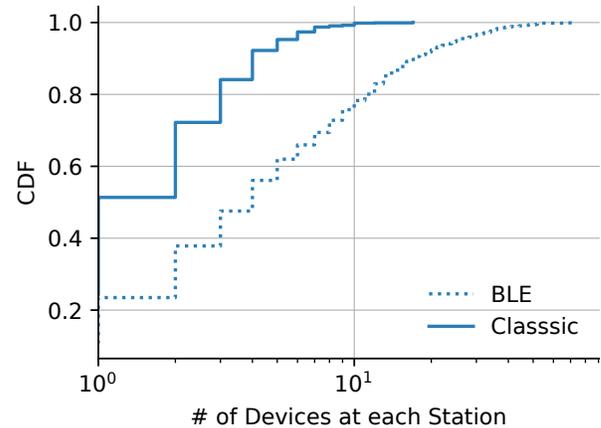


Figure 11: BLE devices are more common than Classic.

Response. BLE devices may be harder to differentiate due to the higher number of devices at each gas station and a lack of distinguishing features. With more sophisticated filtering techniques, it may still be possible to isolate BLE skimmers within this larger set of devices. One possibility is automated RSSI localization to the fuel dispenser location, a possible subject of future research.

5.2 Non-Discoverable Skimmers

The most natural way to evade discovery via Bluetooth would be to put the module in non-discoverable mode. When a Bluetooth device is non-discoverable, it does not respond to normal Bluetooth scans. Instead, it only responds to paging packets specifically addressed to its MAC address.

Cost to attacker. Non-discoverability would make exfiltration more difficult for criminals. One possibility is creating a pre-paired data collection device. However, we have been informed by law enforcement that the individual who installs the skimmer is often independent from the individual responsible for data recovery (called a “mule”). The criminal would not be able to send a mule to recover card data without first delivering them the device. Alternately the criminal could record the MAC address of the skimmer Bluetooth module. This would require careful bookkeeping and the use of tools that support the creation of a non-discoverable connection.

Response. It is still possible to discover a non-discoverable device. For a small set of target address ranges, e.g., $00:06:66$ used by Roving Networks modules, we believe it would be practical to attempt to guess all 16.8 million possible addresses. Prior work has shown that it is possible to discover any non-discoverable device via brute force in 18.64 hours; knowledge of OUI would ideally allow us to reduce this search time [17]. Unfortunately, this requires specialized hardware, rather than an inexpensive Android phone.

5.3 Impersonating Common Benign Devices

Another natural response to Bluetana would be to change the MAC address and name of the device to that of a common benign device, such as a mobile phone or a Bluetooth-enabled car entertainment system. This would make the skimmer appear innocuous to Bluetana.

Cost to attacker. Reprogramming the MAC address on the CSR-based Bluetooth modules, which include the Roving Networks and HC-05 and HC-06 modules, cannot be done using the AT commands used to change device name and pairing. Instead, the skimmer installer would need to re-flash the CSR firmware using a special programming cable. While, in principle, not difficult, it would require an additional degree of sophistication than programming a simple micro-controller development board. The skimmer installer could also change the device name but not the MAC address, say, to one of the known benign devices using the same module, something that is possible to do by issuing AT commands from the micro-controller to the module. While this may cause Bluetana to detect these as a skimmer, signal strength can still be used to identify location of the module.

Response. Because Bluetana collects all Bluetooth data, we can identify skimmers retroactively when we learn of a new MAC address and name used by known skimmers. Thus, if attacks switch to impersonating benign devices, we can update the Bluetana highlighting mechanism to identify those devices as suspicious. This would result in additional inspections, but would still provide significant gain over the state of the art.

5.4 Using Non-Bluetooth Communications

During discussions with law enforcement agencies tasked with identifying skimmers, we were told about skimmers that use GSM modems or WiFi as an alternative to Bluetooth. In the case of WiFi, we believe that the Bluetana methodology will still be effective. GSM poses a more serious challenge for detection.

Cost to attacker. While using GSM would avoid detection using Bluetana, it creates an additional trail of evidence linking the perpetrator to the skimmer. Law enforcement officers could obtain information about the SMS recipient through subpoenas, so receiving the SMS messages on another phone on a US carrier, for example, would be easy to trace. The perpetrator would need to use an SMS service that would not expose his/her identity.

Response. In addition to legal tools available to law enforcement to trace SMS messages, a GSM modem could be detected using a Software-Defined Radio.

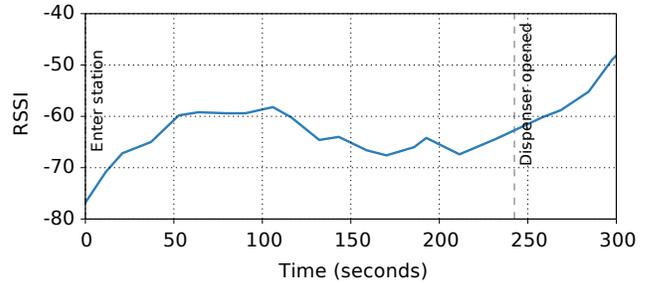


Figure 12: Opening of the gas pump enclosure results in a significant jump in observed Bluetooth signal strength from a skimmer.

5.5 Attacker Bottlenecks

The attacker (skimmer installer) has several practical ways to evade detection using Bluetana. Each of these, however, has an additional cost to the attacker in terms of effort, risk exposure, or expertise. We do not yet have a strong understanding to which of these costs attackers are most sensitive. Indeed, the very low price of stolen credit card numbers, compared to their potential cash out value (Table 1) suggests that the bottleneck in the carding value chain is *not* getting card information but cashing out cards. Thus, while Bluetana may raise the cost for attackers, we do not believe that it will raise it so much as to make fuel dispenser skimming unprofitable.

6 Operational Lessons Learned

While performing the Bluetana study, we learned several lessons about the operational use of Bluetooth scanning for skimmer detection. In this section, we provide an overview of two most important lessons we learned.

6.1 Bluetooth Helps During Inspections

Criminals hide skimmers in the crevices of gas pumps to avoid detection during inspections. We witnessed several instances where investigators were unable to locate skimmers via physical inspection alone. In one incident, Bluetana flagged four devices at a station; however, no skimmers were located. This result led officials more experienced in skimmer recovery to perform a second thorough inspection of the station. These officials located all four skimmers. The evidence provided by Bluetana forced them to continue the inspection, instead of abandoning it and leaving the devices in the field.

Figure 12 demonstrates an instance of how the signal strength measurements helped inspectors determine which pump had a skimmer. When the gas pump's metal door was opened, the signal strength increased significantly, prompting inspectors to look carefully for the skimmer in that pump.

	Group				
	1	2	3	4	5
Skimmers	3	5	6	4	3
Gas stations	2	2	5	4	2
Min. difference in MACs	1	4	9	10	4
Closest MAC distance (in miles)	0	17	59	203	448

Table 5: Several geographically separated skimmers had similar MAC addresses.

6.2 MAC Addresses May Indicate the Source

Network equipment vendors (e.g., Bluetooth module manufacturers) tend to allocate MAC addresses sequentially by production time [34]. Therefore, if two devices have similar MAC addresses, they are likely part of the same batch of devices sold. This information can be used to associate skimmer Bluetooth modules to the same board designer or crew.

We group the skimmers found by Bluetana with the same first 5 bytes of MAC address. Table 5 shows five such groups. We list the difference in MAC address and the geographic distance between the closest MACs in each group. Skimmers in group 1 and 2 were recovered at gas stations in the same county, separated by at most 17 miles. From LE sources, we know that criminals often plant skimmers across multiple stations in a given city/county, and the MAC address data collected indicates this. Groups 3-5 are the most interesting, as the closest MACs in the same group are in stations across different counties. The closest MACs in group 5 are at stations separated by 448 miles. This may seem surprising, but LE informs us that skimmer crews avoid detection by migrating from city to city.

7 Related Work

Skimmer Detection and Prevention. In recent work, Scaife et al. surveyed gas pump skimmer detection and prevention mechanisms [46]. They found that several popular Bluetooth-based skimmer detection applications use a only MAC prefix or device name matching. The results of our study show how the Bluetooth profile of skimmers in these applications can be improved to detect more skimmers, and to flag fewer legitimate devices as skimmers. We also find that Bluetooth-based scanning is an effective way to augment manual gas pump inspections. Scaife et al. also introduced SkimReaper [47], an effective tool for detecting external skimmers. SkimReaper is a credit-card shaped device that an official can swipe in a card reader to detect if the reader has an additional read head: indicating that the reader has an external skimmer attached to it. However, SkimReaper can not detect internal skimmers because they do not add an additional read head. Additionally, the PCI

Security Standards Council have released guidelines for preventing external skimming [41]. Criminals may start using Bluetooth to retrieve card data from external skimmer. If they do, we demonstrate that Bluetooth scanning can augment these existing external skimmer detection and prevention methods.

Bluetooth. Prior work has evaluated the effectiveness of Bluetooth scanning for detecting and localizing Bluetooth devices. They found that Bluetooth signal strength (measured by an Android smartphone) is effective for localizing Bluetooth devices [32, 57]. This work inspired us to use signal strength to detect if a Bluetooth device appears to be installed inside of a gas pump. Previous studies also examined how long it takes to detect a Bluetooth device from stationary observers and moving vehicles. They found that Bluetooth devices are often detected in less time than the Bluetooth standard suggests [39, 43, 24]. This work supports our findings that skimmers are often discovered within the first few seconds of passing by a gas station.

Inventory Attacks. Prior work has demonstrated that user privacy can be violated by inspecting the characteristics of a user’s device [58]. These so called *inventory attacks* have been demonstrated for Bluetooth Low-Energy, RFID, and even web browsers [54, 55, 23]. Our work demonstrates a Bluetooth-based inventory attack against malicious devices, can be used to protect the privacy of consumers.

8 Future Work and Conclusion

As new skimmer detection tools gain popularity, criminals will adapt skimming designs to evade detection. We expect future skimmers will use techniques such those described in Section 5. Similar to Bluetana, future work in this area should emphasize designing easy-to-deploy systems for detecting skimmers, and evaluating their effectiveness with large-scale studies.

Push-back from banks and card issuers has led to wide-scale adoption of EMV in retail PoS systems. However, EMV adoption in gas stations across the U.S. has been slow due to high costs. Therefore, Visa and Mastercard have pushed the EMV adoption deadline for gas stations from 2017 to October 2020 [22]. As gas stations begin migrating to EMV, skimmers targeting EMV will become more common. Future research should focus on the detection of EMV “shimmers” that are gaining in popularity.

Finally, we believe gas pump skimming is the harbinger of an era of attacks using wireless implants. For example, there is an internal Bluetooth-based implant for unlocking door access control systems [14]. Future work should also identify other systems that are vulnerable to using such implants.

In this paper, we presented results of a 19-month long measurement study of Bluetooth scanning as a mechanism to detect internal gas pump skimmers. Our evaluation showed

that Bluetooth characteristics of some internal skimmers can be distinguished from other Bluetooth devices commonly seen at gas stations. We detected, and LE recovered, 64 skimmers at 34 gas stations across four states in the U.S. For 33 of the detected skimmers, Bluetooth was the only source of information that prompted investigators to conduct an inspection. In conclusion, crowdsourced Bluetooth scanning is an effective way to detect Bluetooth-based internal gas pump skimmers.

9 Acknowledgements

We would like to express our appreciation for the local and federal law enforcement agents who introduced us to gas pump skimming and guided us throughout this project. We also like to thank Kevin Allen, and the field investigators at the Arizona Department of Agriculture, Weights and Measures Services Division, for their invaluable help in thoroughly understanding and analyzing the skimming problem in Arizona. We thank the Sacramento County Sheriff's Detectives Sean Smith and Matt Deaux, who are also assigned to the Sacramento Valley Hi-Tech Crimes Task Force, for their help in understanding gas pump skimming in depth. We are also very grateful to the various individuals who drove to gas stations in several states and collected Bluetooth scan data. We also thank our shepherd Joseph Calandrino, and the anonymous reviewers for their insightful feedback and suggestions.

References

- [1] Arizona Department of Agriculture, Weights and Measures Service Division . Data Skimmers in Motor Fuel Dispensers. <https://agriculture.az.gov/sites/default/files/Skimmer%20Presentation%20%28Website%20Edition%29.pdf>, Sept. 2017.
- [2] Nate Seidle . Gas Pump Skimmers . <https://learn.sparkfun.com/tutorials/gas-pump-skimmers>, Sept. 2017.
- [3] Nick Poole . Credit Card Skimmers Evolved: Shimmying . <https://www.sparkfun.com/sparkx/blog/2673>, Apr. 2018.
- [4] Office of Minnesota Attorney General Keith Ellison . ATM and Gas Pump Skimmers . <https://www.ag.state.mn.us/Brochures/pubATMSkimmers.pdf>.
- [5] Rippleshot . State of Card Fraud: 2018. <https://www.aba.com/Products/Endorsed/Documents/Rippleshot-State-of-Card-Fraud.pdf>, 2018.
- [6] United States Sentencing Commission . Guidelines Manual . <https://guidelines.ussc.gov/gl/%C2%A72B1.1>, 2018.
- [7] Affidavit in Support of Criminal Complaints and Arrest Warrants, USA v. Khasanov et al, 1:18cr149. US District Court for the Eastern District of Virginia. <https://www.courtlistener.com/recap/gov.uscourts.vaed.385830/gov.uscourts.vaed.385830.2.0.pdf>, Jan. 2018.
- [8] Appeal from the US District Court for the Eastern District of Oklahoma, USA v. Konstantinov et al, 6:13cr62. United States Court of Appeals for the Tenth Circuit. <https://www.ca10.uscourts.gov/opinions/14/14-7050.pdf>, June 2015.
- [9] Application for Search Warrant, 2:18mj1277. US District Court for the Eastern District of Wisconsin. <https://www.courtlistener.com/recap/gov.uscourts.wied.84529/gov.uscourts.wied.84529.1.0.pdf>, July 2018.
- [10] Arizona Department of Agriculture. Credit Card Skimmers. <https://agriculture.az.gov/weights-measures/fueling/credit-card-skimmers>, Feb. 2019.
- [11] K. Arnold. Florida gas pump thefts rise as credit-card skimmers get more savvy. <https://www.orlandosentinel.com/business/consumer/os-bz-credit-card-skimmers-20181108-story.html>, Nov. 2018.
- [12] ATM Industry Association. Global Fraud and Security Survey - 2017. <https://www.ncr.com/company/blogs/financial/how-much-does-atm-crime-cost>, Jan. 2018.
- [13] H. Bar-El. White Paper: Known Attacks Against Smartcards. Technical report, Discretix Technologies Ltd., 2005.
- [14] M. Bassegio and E. Evenchick. Breaking access controls with BLEKey. <https://www.blackhat.com/docs/us-15/materials/us-15-Evenchick-Breaking-Access-Controls-With-BLEKey-wp.pdf>, Aug. 2015.
- [15] M. Bond, O. Choudary, S. J. Murdoch, S. Skorobogatov, and R. Anderson. Chip and Skim: Cloning EMV Cards with the Pre-play Attack. In *Proc. IEEE Symposium on Security and Privacy*. IEEE, 2014.
- [16] Criminal Complaint, USA v Cristea et al, 4:16cr182. US District Court for the Southern District of Texas. <https://www.courtlistener.com/recap/gov.uscourts.txsd.1357299.1.0.pdf>, Apr. 2016.
- [17] D. Cross, J. Hoeckle, M. Lavine, J. Rubin, and K. Snow. Detecting non-discoverable bluetooth devices. In *International Conference on Critical Infrastructure Protection*, pages 281–293. Springer, 2007.
- [18] The Ultimate Instore Carding by n3d from Darknet. <http://wickybay.com/2017/10/ultimate-instore-carding-n3d-darknet/>.
- [19] DbaseJob. Carding!!! How To Make Your First Money. <https://prvtzone.ws/threads/carding-how-to-make-your-first-money.5052/#post-20315>.
- [20] Tutorial Carding with Dumps. <https://honeymoney24cc.com/cardingwithdumps>.
- [21] CC Dumps Shop. <https://dumps.to/>, Feb. 2019.

- [22] Electronic Transactions Association. ETA Statement on Visa and Mastercard's EMV Liability Shift Date Changes. <https://www.electran.org/eta-statement-on-visa-and-mastercards-emv-liability-shift-date-changes/>, 2016.
- [23] K. Fawaz, K.-H. Kim, and K. G. Shin. Protecting Privacy of BLE Device Users. In *USENIX Security Symposium*, pages 1205–1221, 2016.
- [24] J. Haartsen. Bluetooth—The universal radio interface for ad hoc, wireless connectivity. *Ericsson Review*, 3(1):110–117, 1998.
- [25] Indictment, USA v. Rodriguez et al, 1:17cr417. US District Court for the Northern District of Ohio. <https://www.courtlistener.com/recap/gov.uscourts.ohnd.237118.1.0.pdf>, Oct. 2017.
- [26] Krebs on Security. Skimmers Siphoning Card Data at the Pump. <https://krebsonsecurity.com/2010/07/skimmers-siphoning-card-data-at-the-pump/>, July 2010.
- [27] Krebs on Security. Pro-Grade Point-of-Sale Skimmer. <https://krebsonsecurity.com/2013/02/pro-grade-point-of-sale-skimmer/>, Feb. 2013.
- [28] Krebs on Security. Gang Rigged Pumps With Bluetooth Skimmers. <https://krebsonsecurity.com/2014/01/gang-rigged-pumps-with-bluetooth-skimmers/>, Jan. 2014.
- [29] Krebs on Security. Tracking a Bluetooth Skimmer Gang in Mexico. <https://krebsonsecurity.com/2015/09/tracking-a-bluetooth-skimmer-gang-in-mexico/>, Sept. 2015.
- [30] Krebs on Security. ATM 'Shimmers' Target Chip-Based Cards. <https://krebsonsecurity.com/2017/01/atm-shimmers-target-chip-based-cards/>, Jan. 2017.
- [31] Legitshop. Trusted Dumps with PIN. <https://legitshop.org/>, Feb. 2019.
- [32] S. Liu, Y. Jiang, and A. Striegel. Face-to-face proximity estimation using bluetooth on smartphones. *IEEE Transactions on Mobile Computing (TMC)*, 13(4):811–823, 2014.
- [33] D. MacGuill. Can a Mobile Phone's Bluetooth Sensor Be Used to Detect Card Skimmers? <https://www.snopes.com/fact-check/bluetooth-gas-pump-skimmers/>, 2019.
- [34] J. Martin, E. Rye, and R. Beverly. Decomposition of MAC address structure for granular device inference. In *Proc. Annual Computer Security Applications Conference (ACSAC)*, 2016.
- [35] Meccadumps. Buy Dumps CVV online Fullz Verified seller. <https://meccadumps.net/>, Feb. 2019.
- [36] Memorandum and Order, USA v. Hristov et al, 1:10cr10056. US District Court for the District of Massachusetts. <https://www.courtlistener.com/recap/gov.uscourts.mad.127405/gov.uscourts.mad.127405.62.0.pdf>, Apr. 2011.
- [37] Mettler-Toledo. BC Shipping Scale Service Manual. <https://thescalestore.com/manuals/Mettler-Toledo-BC-User-Manual-1.pdf>, Aug. 2015.
- [38] MH Corbin Highway Information Systems. Surface Scan. [http://mhcorbin.com/Portals/0/MH%20Corbin%20Surface%20Scan%20User%20Manual%20v1.1%20\(002\)%20new%20cover.pdf](http://mhcorbin.com/Portals/0/MH%20Corbin%20Surface%20Scan%20User%20Manual%20v1.1%20(002)%20new%20cover.pdf), Jan. 2018.
- [39] P. Murphy, E. Welsh, and P. Frantz. Using bluetooth for short-term ad-hoc connections between moving vehicles: A feasibility study. In *IEEE Vehicular Technology Conference (VTC)*, volume 1, 2002.
- [40] Everything you need to know about instore carding. <http://wickybay.com/2017/11/everything-need-know-instore-carding/>, Nov. 2017.
- [41] PCI Security Standards Council. Skimming Prevention: Overview of Best Practices for Merchants. https://www.pcisecuritystandards.org/documents/Skimming_Prevention_At-a-Glance_Sept2014.pdf, Sept. 2014.
- [42] PCI Security Standards Council. PCI DSS Quick Reference Guide. https://www.pcisecuritystandards.org/documents/PCI_DSS-QRG-v3_2_1.pdf, 2018.
- [43] B. S. Peterson, R. O. Baldwin, and J. P. Kharoufeh. Bluetooth inquiry time characterization and selection. *IEEE Transactions on Mobile Computing (TMC)*, 5, 2006.
- [44] PRTSHIP. DUMPS. <https://prtship.com/forums/dumps.6/>.
- [45] Santander Bank. What is my debit card spending/withdrawal limit? https://customerservice.santanderbank.com/app/answers/detail/a_id/3713/kw/atm%20withdraw/r_id/102441.
- [46] N. Scaife, J. Bowers, C. Peeters, G. Hernandez, I. N. Sherman, P. Traynor, and L. Anthony. Kiss from a rogue: Evaluating detectability of pay-at-the-pump card skimmers. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1208–1222, Los Alamitos, CA, USA, may 2019. IEEE Computer Society.
- [47] N. Scaife, C. Peeters, and P. Traynor. Fear the Reaper: Characterization and Fast Detection of Card Skimmers. In *Proc. USENIX Security*, 2018.
- [48] Scientific Working Group on Digital Evidence. Best Practices for Examining Magnetic Card Readers. <https://www.swgde.org/documents/Current%20Documents/SWGDE%20Best%20Practices%20for%20Examining%20Magnetic%20Card%20Readers>.
- [49] Sell CVV (CC). <https://sellcvvdumps.shop/>.
- [50] Sentencing Memorandum of the United States, USA v. Aqel, 2:14cr270. US District Court for the Southern District of Ohio. <https://www.courtlistener.com/recap/gov.uscourts.ohsd.178108/gov.uscourts.ohsd.178108.47.0.pdf>, Nov. 2015.
- [51] Skim Plus (Bluetooth Skimmer Detection). <https://play.google.com/store/apps/details?id=com.rs.skimplus.beta>, 2018.

- [52] Teletrac. Teletrac Drive User Guide. http://community.teletrac.com/teletrac.com/assets/2014-04-23_android%20tablet%20user%20guide.pdf, Jan. 2014.
- [53] The Newnan Times-Herald. Armenian skimmer leader pleads guilty. <http://times-herald.com/news/2015/06/armenian-skimmer-leader-pleads-guilty>, July 2017.
- [54] T. van Deursen. 50 ways to break RFID privacy. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 192–205. Springer, 2010.
- [55] A. Vastel, P. Laperdrix, W. Rudametkin, and R. Rouvoy. Fp-Scanner: The Privacy Implications of Browser Fingerprint Inconsistencies. In *Proc. USENIX Security*. USENIX Association, 2018.
- [56] VICE. Gangs on the Dark Web: Credit Card Scammers. <https://www.youtube.com/watch?v=jT-jmq8KBw0>, June 2018.
- [57] Y. Wang, X. Yang, Y. Zhao, Y. Liu, and L. Cuthbert. Bluetooth positioning using RSSI and triangulation methods. In *Consumer Communications and Networking Conference (CCNC), 2013 IEEE*, pages 837–842. IEEE, 2013.
- [58] J. H. Ziegeldorf, O. G. Morchon, and K. Wehrle. Privacy in the internet of things: threats and challenges. *Security and Communication Networks*, 7(12):2728–2742, 2014.

A Court Cases

The appendix contains excerpts from various public court documents related to cases of credit card skimming. These excerpts provide anecdotal data about the monetary impact of the skimmer problem.

A.1 Cashout Value

USA v. Hristov et al [36]

"... Bank of America suffered a loss of \$33,000 with 36 compromised customer accounts. Citizens Bank suffered a loss of \$91,580 with 74 compromised customer accounts ..."

USA v. Cristea et al [16]

"... Altogether, on February 21, 2016, FBI surveillance observed Cristea, Co-conspirator #1, and Co-conspirator #2 go to approximately 12 different locations, where, according to CardTronic's records, they withdrew at least \$7,000 from at least 18 First National Bank accounts ..."

USA v. Khasanov et al [7]

"... USPS agents thereafter conducted record checks on the purchased USPS money orders and discovered that 10 of the 57

money orders had been purchased with 5 payment numbers issued by Citibank ..."

Date	Location of USPS	Amount
Aug 4 2017	Waldorf, MD	\$2,904.80
Aug 7 2017	Washington, DC	\$1,492.80
Aug 7 2017	McLean, VA	\$1,400.00
Aug 7 2017	Washington, DC	\$1,803.20
Aug 7 2017	Hyattsville, MD	\$792.05

USA v. Aqel [50]

"... the Probation Officer also notes that the actual loss to victims was \$8,327.58. Id. Similarly, the Probation Officer notes that while Mr. Aqel possessed 120 stolen credit card numbers, only 23 of those numbers were used to make purchases ..."

USA v. Rodriguez et al [25]

"... Between on or about July 7, 2016, and on or about July 20, 2016, Defendant ... attempted to conduct approximately 133 retail transactions totaling in excess of \$27,000 ... using approximately 90 counterfeit access devices re-encoded with credit/debit account information that were obtained by a skimming device placed on the point of sale terminal of a gas pump ..."

Application for Search Warrant, 2:18mj1277[9]

"... On April 14, 2016, a man (later identified as Estrada) used a fraudulent Visa credit card and a fraudulent MasterCard to purchase two \$300.00 gift cards from the Kohl's store ..."

USA v. Konstantinov et al [8]

"... In total, the defendants compromised approximately 524 debit card accounts and made approximately 779 fraudulent withdrawals, totaling \$348,376.80 ..."

A.2 Credit/Debit cards per skimmer per day

Application for Search Warrant, 2:18mj1277 [9]

"... On September 9, 2016, an employee at Jilly's Mobil ... reported to Detective Craig Meyer that he had found what appeared to be a skimmer on pump #8 ... Detective Meyer downloaded and exported the data stored on the skimmer taken from Jilly's Mobil pump #8. The results showed data for 221 victim credit card accounts ...
... Detective Meyer reviewed the video surveillance footage for Jilly's Mobil from September 1, 2016. At 1:38 PM on September 1st, a red Ford Explorer drove to pump 8. The Ford Explorer was positioned in a manner whereby the opened passenger door blocked the view of the gas pump by the store employee inside the Jilly's Mobil ..."