Securing the Wireless Emergency Alerts System

By Jihoon Lee, Gyuhong Lee, Jinsung Lee, Youngbin Im, Max Hollingsworth, Eric Wustrow, Dirk Grunwald, and Sangtae Ha

Abstract

Modern cell phones are required to receive and display alerts via the Wireless Emergency Alert (WEA) program, under the mandate of the Warning, Alert, and Response Act of 2006. These alerts include AMBER alerts, severe weather alerts, and (unblockable) Presidential Alerts, intended to inform the public of imminent threats. Recently, a test Presidential Alert was sent to all capable phones in the U.S., prompting concerns about how the underlying WEA protocol could be misused or attacked. In this paper, we investigate the details of this system and develop and demonstrate the first practical spoofing attack on Presidential Alerts, using commercially available hardware and modified open source software. Our attack can be performed using a commercially available software-defined radio, and our modifications to the open source software libraries. We find that with only four malicious portable base stations of a single Watt of transmit power each, almost all of a 50,000-seat stadium can be attacked with a 90% success rate. The real impact of such an attack would, of course, depend on the density of cellphones in range; fake alerts in crowded cities or stadiums could potentially result in cascades of panic. Fixing this problem will require a large collaborative effort between carriers, government stakeholders, and cellphone manufacturers. To seed this effort, we also propose three mitigation solutions to address this threat.

1. INTRODUCTION

The Wireless Emergency Alerts (WEA) program is a government mandated service in commercialized cellular networks in the U.S. WEA was established by the Federal Communications Commission (FCC) in response to the Warning, Alert, and Response Act of 2006 to allow wireless cellular service providers to send geographically targeted emergency alerts to their subscribers. The Federal Emergency Management Agency (FEMA) is responsible for the implementation and administration of a major component of WEA.

This system can send three types of alerts: **Presidential Alerts** issued by the president to all of the United States; **Imminent Threat Alerts** involving serious threats to life and property, often related to severe weather; and **AMBER Alerts** regarding missing or abducted children. Considering the number of cellphone users and the nationwide coverage of cellular networks, WEA over Long-Term Evolution (LTE) was a natural step to enhance public safety *immediately* and *effectively*. In fact, recent rapidly moving fires have caused emergency services to consider using WEA instead of relying on opt-in alerting systems.¹⁶

A handful of widely publicized events has led to public scrutiny over the potential misuse of the alert system. On

January 13, 2018, there was a geographically targeted alert issued in Hawaii. The message, warning of an inbound missile, is shown in Figure 1b. Although caused by human error, the impact on the residents of Hawaii was huge, as it led to panic and disruption throughout the state.²⁰ This event was followed on October 3, 2018, by the first national test of a mandatory Presidential Alert. The alert, captured in Figure 1a, was sent to all capable phones in the U.S.¹⁹

These recent high-profile alerts have prompted us to assess the realizability and impact of an alert spoofing attack. In this paper, we demonstrate how to launch a Presidential Alert-spoofing attack and evaluate its effectiveness with respect to attack coverage and success rate.

To answer this question, we start by looking into the alert delivery method used by WEA. WEA sends alerts via the commercial mobile alert service (CMAS), which is the underlying delivery technology standardized by the 3rd Generation Partnership Project (3GPP). These alerts are delivered via the LTE downlink within broadcast messages, called System Information Block (SIB) messages. A celltower (referred to as eNodeB) broadcasts the SIB to every cell phone (referred to as user equipment or UE) that is tuned to the control channels of that eNodeB. A UE obtains necessary access information, such as the network identifier and access restrictions, from SIB messages, and uses it for the eNodeB selection procedure. Among the 26 different types of SIB messages, SIB12 contains the CMAS notification, which delivers the aforementioned alert messages to the UEs.

The eNodeB broadcasts SIB messages to the UE, independently from the mutual authentication procedure that

Figure 1. Snapshots of real WEA messages received by cellphones: (a) the first national test of the Presidential Alert performed on October 3, 2018 in the U.S., and (b) a false alert sent in Hawaii on January 13, 2018.

(a) Presidential alert

(b) Imminent threat alert

The original version of this paper is entitled “This is Your President Speaking: Spoofing Alerts in 4G LTE Networks” and was published in *Proceedings of the 17th ACM International Conference on Mobile Systems, Applications and Services*, 2019.

eventually occurs between them. Thus, all SIBs, such as CMAS, are intrinsically *vulnerable* to spoofing from a malicious eNodeB. More importantly, even if the UE has completed its authentication and securely communicates with a trusted eNodeB, the UE is still exposed to the security threat caused by the broadcasts from other, possibly malicious, eNodeBs. This is because the UE periodically gathers SIB information from neighboring eNodeBs for potential eNodeB (re)selection and handover.

We found via both experiment and simulation that a 90% success rate can be reached in 4435 m² of a 16,859 m² building using a single malicious eNodeB of 0.1 Watt power, whereas in an outdoor stadium, 49,300 seats among the total 50,000 are hit with an attack, which itself has a 90% success rate using four malicious eNodeBs of 1 Watt power.

In summary, we make following major contributions:

- We identify security vulnerabilities of the WEA system and explain the detailed underlying mechanism stipulated by the LTE standard. We find that the CMAS spoofing attack is easy to perform but is challenging to defend in practice.
- We present our threat analysis on the CMAS spoofing attack and implement an effective attack system using commercial off-the-shelf (COTS) software-defined radio (SDR) hardware and open-source LTE software.
- We evaluate our attack system using both SDR-based hardware prototype and measurement-based simulation. As one of the striking results, we demonstrate that four SDR-based malicious eNodeBs at 1 Watt of power can propagate their signal to 49,300 of the whole 50,000-seat football stadium. Of the 49,300 seats affected, 90% will receive the CMAS message.
- We present possible solutions to prevent such a spoofing attack with a thorough analysis and feasibility test, which can open the door toward collaborative efforts between cellular operators, government stakeholders, and phone manufacturers.

1.1. Responsible disclosure

In January 2019, before public release, we disclosed the discoveries and technical details of this alert spoofing attack to various pertinent parties. These parties include the government and standardization organizations FEMA, FCC, DHS, NIST, 3GPP, and GSMA; the cellular network service

providers AT&T, Verizon, T-Mobile, Sprint, and U.S. Cellular; and the manufacturers Apple, Google, and Samsung.

2. SECURITY THREATS

The 3GPP standardization body began a project in 2006 to define the requirements of CMAS to deliver WEA messages in the LTE network, and the LTE CMAS network architecture is illustrated in Figure 2. The resulting technical specification, initially released in 2009, describes the general criteria for the delivery of alerts, message formats, and functionality of CMAS-capable UEs.² During an emergency, authorized public safety officials send alert messages to Federal Alert Gateways. The participating mobile service providers then broadcast the alert to the UEs, who will automatically receive the alert if they are located in or travel to the targeted geographic area. The cell broadcast center (CBC) is part of the service provider's core network and is connected to the Mobility Management Entity (MME), which maintains the location information of the UEs attached to the network.³ The eNodeB is the final step in communicating the alert to the UEs over the air.

UEs may choose to turn off the notification of imminent threat alerts and AMBER alerts among the three types of emergency alerts (i.e., presidential alerts, imminent threat alerts, and AMBER alerts). However, the 3GPP mandated that the reception of Presidential Alerts is obligatory. Thus, cell phones have no option to disable Presidential Alerts, as seen in Figure 3. Because it cannot be disabled, this paper focuses on spoofing Presidential Alerts with the injection of a fake CMAS message over the air from a rogue eNodeB.

2.1. Identifying the vulnerability

An eNodeB broadcasts LTE system information through the Master Information Block (MIB) and SIB. Specifically, when a LTE searches for an eNodeB, it searches for the eNodeB's physical cell identifier (PCI) within a dedicated synchronization channel specified by the LTE standard.⁵ After finding the PCI, the LTE unscrambles the MIB, which contains essential information such as the system bandwidth, system frame number (SFN), and the antenna configuration, to decode the SIB Type 1 message (SIB1). There are several SIB messages but only SIB1 has a fixed periodicity of 80

Figure 2. LTE CMAS network architecture.

Figure 3. Government alert settings in mobile phones: (a) Android and (b) Apple's iPhones. Although AMBER and emergency alerts can be manually disabled, users cannot disable or block Presidential Alerts from being received or displayed.

msec. Other SIB messages are dynamically scheduled by the eNodeB, and the scheduling information for other SIBs is encoded in the periodic SIB1.

3GPP specifies that the broadcast of CMAS messages is over the air through SIB12.⁶ However, unlike point-to-point messages in LTE, broadcasts of SIB messages are not protected by mutual cryptographic authentication or confidentiality, because the SIB contains essential information the UEs use to access the network before any session keys have been established. Once a CMAS message has been received, there is no verification method for the message content. If an attacker can imitate eNodeB behavior closely enough to broadcast false CMAS messages, the UE will display them.

A UE's vulnerability to a fake CMAS alert depends on whether it is in an *active* or *idle* state, illustrated in Figure 4. To affect the most UEs, an attacker must consider different approaches for each state. Here we discuss idle UEs and active UEs separately:

Idle mode UEs. Reference Signal Received Power (RSRP) is the power of an eNodeB-specific reference signal recognized by the UE, typically used to make an eNodeB selection and handover decision. Usually, whenever a UE in idle mode performs eNodeB selection (or reselection), it will associate with the eNodeB having the highest RSRP. If the RSRP of a malicious eNodeB is the strongest, the UE decodes the SIBs transmitted by the malicious eNodeB. The attacker does not need to have any user information (such as security keys), which would be stored in the network operator's database. Without having such user information, the UE will eventually reject the authentication process with the malicious eNodeB. However, it can receive a CMAS message transmitted by the malicious eNodeB during this process.

Active mode UEs. When a UE is in active mode, it securely communicates with the serving eNodeB. If it finds another eNodeB with a higher power level than the existing serving eNodeB, a handover procedure can be triggered. The serving eNodeB then makes a handover decision based on the received measurement report. However, if the serving MME

does not identify the target eNodeB, the handover will eventually fail. Therefore, even if caused by a malicious eNodeB, the handover procedure does not make a UE vulnerable to the CMAS spoofing attack. As a consequence, the attacker first needs to disconnect the UE from its serving eNodeB. After the UE is released from the serving eNodeB, it will immediately try to attach to the strongest eNodeB. After that, it can be attacked in the same way as idle mode UEs described in the section above. One way to disconnect the active UE from its serving eNodeB is to incur Radio Link Failures (RLFs) by jamming LTE signals.¹⁵ Simply, without any special jamming technique, a malicious eNodeB can jam the communication between a UE and its serving eNodeB by merely transmitting at a much higher power than the serving eNodeB.

2.2. CMAS reception and trustworthiness

We have identified three possible cases that determine whether the CMAS is received and is trustworthy in Table 1. Each case depends on where the UE is currently in the idle/active life cycle, illustrated in Figure 4.

Simply put, if a UE is not listening to frequency channels on which the eNodeB is transmitting the CMAS message, the CMAS message will not be received by the UE. This is illustrated as the blue portion in Figure 4. It may seem obvious, but a necessary condition for the UE to receive a CMAS message is that it needs to be tuned to the synchronization channels of the eNodeB that is transmitting the CMAS message.

Secure CMAS. In the green area of Figure 4, the UE attaches to an eNodeB and is safely in the active state. To do this, the UE must be equipped with a valid Service Identity Module (SIM) card that is registered to the operator's network. Case 1 is the general scenario for phones receiving standard service from their provider. Because mutual authentication between the UE and the network has been successfully made, the UE can trust that the eNodeB is not malicious. The CMAS reception is successful as we would expect, and we know that this CMAS message is trustworthy.

Unsecured CMAS. In the red area of Figure 4, the UE is failing or has already failed to attach when the eNodeB transmits the CMAS message. The UE will still receive the CMAS message; this is the crux of the vulnerability. To demonstrate this, we deleted the SIM information from the Evolved Packet Core (EPC) so that the user authentication would be unsuccessful. The UE is now in the unsecured range between the idle and active states due to the authentication failure. Even though the UE fails to reach the active state, we observe that the CMAS message is still successfully received. This is because once the UE completes decoding the CMAS message in SIB12, it delivers the contents to the application layer to be shown to the user. Surprisingly, this is possible even after

Figure 4. The Idle/Active life cycle of a UE. The state of the UE continues counterclockwise around the chart. CMAS spoofing is possible although the UE performs an eNodeB search, prior to successful authentication with a trusted eNodeB.

Table 1. Cases for CMAS reception and trustworthiness.

Case	SIM equipped	Auth. success	CMAS reception	Trustworthy
1	Yes	Yes	Yes	Yes
2	Yes	No	Yes	No
3	No	No	Yes	No

the authentication process has finally failed. Case 2 can lead the potential threat that *any malicious eNodeB can deliver fake CMAS messages although the UE is in between the eNodeB search and authentication procedures*. Finally, in Case 3, the UE roams to an eNodeB, which sends a CMAS message. To demonstrate this, we removed the SIM card from the UE. No authentication is possible, but the UE can make emergency calls such as 911. Even in this situation, we verified that the UE still receives the CMAS message, which is potentially malicious.

As shown in Cases 2 and 3, CMAS spoofing can be done although the UE performs an eNodeB search before successful authentication with a trusted eNodeB. These results are verified using $1 \times$ JLG20 COTS LTE small cell (no modification), $1 \times$ open-source NextEPC (modified with the CBC),¹⁷ and nine different commercial LTE phones (Apple iPhone 8, X, and XS; Google Pixel 1; Huawei Nexus 6P; Motorola G5 Plus and G6; Samsung Galaxy S7 Edge and S8). Considering that the majority of UEs in cellular networks are in the idle state¹⁰ and UEs often transition from the active to idle state due to an inactivity timer (around 10 s^{13}), *almost all UEs are susceptible to this attack*.

3. PROOF-OF-CONCEPT ATTACKS

In this section, we present the details of our *Presidential Alert Spoofer* system and describe how it works. Our system can be built with either an SDR device or a COTS eNodeB, and the list of hardware and software systems we used is summarized in Table 2.

Attack preparation. Our Presidential Alert Spoofer must first identify the existing eNodeBs in a given licensed frequency band. Each eNodeB can be uniquely identified at a given geographical position by the pair of “E-UTRA Absolute Radio Frequency Channel Number (EARFCN)” and “Physical Cell ID (PCI).” For each EARFCN, our Spoofer finds the eNodeB, and associated PCI, of which the RSRP is the strongest. Once the existing eNodeBs are listed, the Public Land Mobile Network (PLMN) information of each eNodeB is collected. Every LTE network has its PLMN, a three-digit country code, and two or three digits to identify the provider. The PLMN is periodically broadcast by the eNodeB in the SIB1 message, making it possible to collect all of the observable PLMNs within the receiving range passively. To launch an attack, our Presidential Alert Spoofer uses the same PLMN as an existing eNodeB such that the UEs will select our Spoofer during an eNodeB search.

Table 2. HW and SW systems used for implementation.

System	Hardware	Software
Attack preparation	BladeRF 2.0 (\$500) USRP B210 (\$1300) Laptop (< \$1000)	OWL ⁸ (modified)
SDR-based Spoofer	BladeRF 2.0 (\$500) USRP B210 (\$1300) Laptop (< \$1000)	srsLTE ¹² (modified)
COTS eNodeB-based Spoofer	JLG20 (FDD) JLT621 (TDD) Laptop (< \$1000)	NextEPC ¹⁷ (modified)

Attack execution with an SDR device. We implemented the Spoofer using a USRP B210 and BladeRF to attack Frequency Division Duplex (FDD) systems. With an SDR, we can change the transmission frequency easily to target every cellular band. We added SIB12 support to the open-source eNodeB software¹² and could transmit a CMAS message every 160 msec.

Attack execution with a COTS eNodeB. We use a COTS eNodeB (Juni JLT-621) to target Time Division Duplex (TDD) systems. Our modification of NextEPC provides an interface to inject a user-defined Presidential Alert that broadcasts each second. With this configuration, a victim UE may receive the SIB12 every second from the COTS eNodeB. Any commercial LTE FDD/TDD eNodeB hardware can perform this attack, which may play a key role if an attacker wants to control multiple malicious eNodeBs in a coordinated manner.

Attack verification. In our lab environment, we verified that the fake Presidential Alert sent by our SDR-based

Figure 5. The Presidential Alert Spoofer scans for an eNodeB, gathers operator information, and sends a fake Presidential Alert to both idle and active UEs. The UEs may be FDD or TDD. This setup consists of one SDR device, one COTS LTE eNodeB, and two laptops.

Figure 6. Receiving multiple fake Presidential Alerts using a Samsung Galaxy S8(left) and an Apple iPhone X(right).

Spoofers were successfully shown in the FDD phones of AT&T, T-Mobile, and Verizon. With a TDD Sprint phone, we verified that our COTS eNodeB-based Spoofer also works successfully. All the experiments are carried out with proper RF shielding.

Affected devices and implications. From discussions of the SIB12 vulnerability in §2.1, it became clear that the lack of authentication was a design choice by 3GPP, rather than an oversight. This design provides the best possible coverage for legitimate emergency alerts, but the trade-off leaves every phone vulnerable to spoofed alerts. Consequently, all modem chipsets that fully comply with the 3GPP standards show the same behavior: the fake Presidential Alert is received without authentication. Once the LTE modem of the UE receives the fake alert, the operating system will display the alert to the user. Because our attack verification tests included many Android and iOS phones, we conclude that most (presumably all) LTE phones will be affected by the attack, regardless of the phone's vendor or model. Moreover, much of the LTE public warning system is inherited from 2G/3G and continues in 5G; a similar attack is also possible in 5G.

4. EVALUATION

Figure 7 illustrates our experimental testbed setup, which consists of an EPC and eNodeB for a conventional LTE system, a malicious eNodeB for spoofing, and cell phones for victim UEs. A signal attenuator receives the broadcast signals from two sources and delivers the combined signal to a LTE in a shielded box. We built an LTE test network with an EPC and eNodeB, named SecureNet, which assumes the role of the user's original network. On the other hand, the malicious eNodeB, part of the Presidential Alert Spoofer, is installed solely without any LTE core support. By using the signal attenuator, the signal power received at the LTE can be precisely controlled for various practical scenarios.

4.1. Success rate

Let α be the RSRP difference between the SecureNet eNodeB and Presidential Alert Spoofer for an idle UE (i.e., $\alpha = RSRP_{SecureNet} - RSRP_{Spoofer}$) and β be the RSRP difference for an active UE. Then we evaluate the Presidential Alert Spoofer's success rate as a function of α (or β). We first attach the UE to SecureNet. For the idle UE case, we wait for

Figure 7. The testbed setup for evaluating the attack success rate. The transmission power levels of the SecureNet eNodeB and the Presidential Alert Spoofer can be controlled independently.

the UE to enter the idle mode due to inactivity. The Spoofer broadcasts each new Presidential Alert message, so we can determine whether each Presidential Alert is successfully received and at what power configuration of α or β . We conducted 20 experimental trials for each value of α (or β) ranging from 0 to -25 dB.

The Spoofer may elect to use a different PCI than that of the serving eNodeB, appearing to be a new eNodeB. Or, the Spoofer may use the same PCI, looking to be the existing eNodeB and interfering with the existing eNodeB's PHY-layer control channel information.²² This decision has different impacts on the performance of the spoofing attack, depending on the UE state (*idle* or *active*).

Figure 8 shows the empirical cumulative distribution function (CDF) of successful receptions of fake alerts as a function of α for idle UEs. When the Spoofer uses a different PCI and the received signal strength from the Spoofer is higher than that from SecureNet ($\alpha < 0$), the idle UE will consider the Spoofer as a new serving eNodeB. Our experimental results verify this expectation; 50% of idle UEs can receive a fake message even at $\alpha = -1$, and more than 90% of idle UEs can receive a fake message when $\alpha \leq -6$.

However, if the same PCI is used, the attack performance is significantly degraded. Because the PCI is used to generate cell-specific reference signals,⁵ using the same PCI value will cause channel estimation errors at the UE due to collisions from the two transmitters. This, in turn, leads to more decoding errors when receiving the SIBs. As a result, using the same PCI requires much higher attack power as no UE is affected when α is greater than -12 dB. With $\alpha \leq -17$, 90% of idle UEs can still be attacked.

Figure 9 shows the CDF of successful fake message receptions as a function of β (i.e., forcing disconnect) for active UEs. When the Spoofer uses a different PCI, and the received signal strength from the Spoofer is higher than that from the SecureNet eNodeB, the active UE will start to consider the Spoofer as a target eNodeB for a handover, as described in §2.2. Because SecureNet does not identify the Spoofer, a handover cannot be performed. Instead, we observed an RLF would occur when $\beta \leq -10$, which eventually leads to the

Figure 8. The CDF as a function of α for only idle UEs. Because eNodeB reselection happens when idle UEs wake up, the spoofing attack performs better when using a different PCI.

reception of a fake alert. About 90% of active UEs can receive a fake message when $\beta \leq -20$, assuming that a different PCI value is used for the Spoofer. Unlike the idle UE case, using the same PCI value results in higher decoding errors (and more RLFs) at a receiver. Thus, it shows better attack performance; 90% of receptions are successful with $\beta \leq -13$.

4.2. Practical scenarios: indoor and outdoor

As we do not use the Spoofer outside of a shield box, we cannot directly measure its effect on a large number of people. To evaluate the attack coverage according to its success rate, we use actual RSRP measurements in indoor and outdoor environments.

Indoor attack. We placed our malicious eNodeB inside a campus building and measured the RSRP of a dummy LTE signal (containing no CMAS message) in the EBS band with 0.1 Watt transmit power. We also measured the RSRP of a nearby AT&T eNodeB, as shown in Figure 10a. The RSRP does not attenuate consistently due to various obstacles, but generally, the RSRP tends to decrease as the distance from

Figure 9. The CDF as a function of β for only active UEs. Using the same PCI leads to more decoding errors observed by the UE. This results in a slightly more effective attack.

Figure 10. The indoor attack simulation: (a) the satellite image of the Engineering Center at the University of Colorado Boulder shows the nearest AT&T eNodeB. The graph shows the indoor RSRP distribution of that eNodeB. (b) The attack coverage for idle and active UEs are shown when a 1×0.1 Watt malicious eNodeB is used.

the AT&T eNodeB increases. We compared the two RSRPs throughout the building and indicated the attack coverage using measurements obtained from §4.1, as depicted in Figure 10b. As a result, in a building with a total area of about 16,859 m², for idle UEs, the coverage for a 90% success rate was about 4435 m², whereas for active UEs, the coverage for a 90% success rate was about 2955 m².

Outdoor attack. Without access to outdoor LTE equipment, we simulate the RSRPs of the spoofing eNodeB and the AT&T eNodeB with the NS-3 v3.29 network simulator.¹⁸ For the scenario, we assume a football game where a large number of people are gathered in a restricted region. A group of attackers sends fake alerts to the spectators inside the football stadium. We measured the RSRP of an actual AT&T eNodeB around the perimeter of our campus' football stadium, as shown in Figure 11. We used the simulator to estimate the RSRPs at the centers of each section in the stadium (Figure 11a). We simulated the spoofers in four corners around the stadium, near but still outside of the ticketed area. Figure 11b shows which malicious eNodeB with a 1 Watt transmit power attacked each section. We observe that all sections, except one, are attacked by the malicious eNodeBs. This means that 49,300 among the total 50,000 seats will be hit with the attack, which itself has a 90% success rate, given that all UEs are in the idle state.

5. MITIGATION SOLUTIONS

Defending against CMAS spoofing attacks requires careful consideration of several challenges. First, updates to the CMAS architecture could require expensive changes by cell phone manufacturers, operating system developers, government bodies, and cellular carriers. Coordinating such an effort would be difficult due to the fragmented nature of the network. Furthermore, updates must still support outdated devices, both on the user (UE) and infrastructure (eNodeB) side, as it could take years to replace old equipment. Also, any comprehensive defense must consider the trade-off between security and availability: if users cannot receive valid alerts due to sophisticated protections, it may be more

Figure 11. The outdoor attack simulation: (a) the satellite image of Folsom Field at the University of Colorado Boulder shows the location of the AT&T eNodeB. The stadium graph represents the RSRP distribution of the eNodeB measured at the center of each section, (b) When 4×1 Watt malicious eNodeBs are located outside the four corners of the stadium, the simulated attack coverage hits all but one section. This means that 49,300 among the total 50,000 seats are hit with the attack, which itself has a 90% success rate.

hazardous than the case if we continued to use the existing (but vulnerable) system.

With these challenges in mind, we propose three mitigation solutions: first, a client-side software solution ignoring unsecured CMAS alerts; second, a network-aware solution attempting to detect false alerts by modeling characteristics of legitimate eNodeBs; and third, adding digital signatures to alerts.

5.1. Client-driven approach

A client-driven approach should provide an ability for a UE to decide whether a received CMAS message is trustworthy. It requires the information from LTE's control plane, which is responsible for essential operations such as network attaches, security control, authentication, setting up of bearers, and mobility management. To mitigate the CMAS spoofing attack, we utilize Radio Resource Control (RRC) and Non-Access Stratum (NAS) layer information from the LTE control plane. We can check whether the UE has a valid connection or not from the RRC control information and the UE's state transition with MME from the NAS control information.

Monitoring RRC and NAS on a UE is currently tricky because LTE control plane protocols are handled by the LTE baseband chipset and firmware so that accessing such information through the existing Operating System (e.g., Android, iOS) is not fully supported. In our implementation, we installed a cellular debugging tool on Android to retrieve the state information of RRC and NAS.¹⁴

Figure 12 shows the RRC and NAS state transition in a standard scenario where the UE receives a *Secure* CMAS message from a legitimate eNodeB. When it receives an *Unsecure* CMAS message, we will see a different state transition. For instance, when a UE is in active mode, the attack starts with a sudden radio link failure, as we explained in §2. It incurs the RRC state change from "CONNECTED" to "IDLE," and the state goes back to "CONNECTED" when a CMAS is received. After that, the NAS state will soon change into "EMM-REGISTERED.NO-CELL-AVAILABLE."

Figure 12. UE state transition for Secure CMAS reception.

Figure 13. The client-driven approach evaluates the security of the broadcast radio channel by monitoring UE's RRC and NAS state transitions when a CMAS message is received. a fake CMAS message with a warning, as shown in Figure 13.

As a result, we propose a spoofing detection algorithm as a client-driven approach. First, it needs to have the ability to access a short history of RRC and NAS state transition. Then, whenever a CMAS message is received, it should be checked by our algorithm before displaying it to a user. The algorithm finds any suspicious activity by evaluating RRC and NAS state transition, assuming that unsecured connections may deliver fake broadcast messages. Finally, it shows:

5.2. Network-aware approach

A network-aware approach can leverage the received signal strength (RSS) at the UE to determine if the eNodeB from which the UE received the CMAS message is within a feasible distance. Using a widely used path-loss model,¹¹ we can estimate the distance to the eNodeB using the RSS value. Then compare this with the location provided by an Internet database⁹ to determine whether the alert could have come from a trusted eNodeB.

The performance of this technique could be greatly improved by applying a machine learning (ML) as shown in Figure 14. In our design, we train legitimate cells using basic cell information, neighbor relations, and signal quality measurements associated with the location. Such information may be collected and shared by network operators or crowdsourcing.²¹ In our prototype, a UE retrieves an ML model associated with its serving and surrounding cells of its location to classify the validity of the attached eNodeB upon reception of a CMAS message.

5.3. Digital signature approach

We also consider digitally signing SIB12 messages to prevent spoofed messages, as discussed by 3GPP.¹ Although it is conceptually simple, adding signatures is difficult because operators and devices must agree on the key or keys that will be used to sign and validate messages.

For key management, we leverage suggestions from 3GPP discussions,¹ which suggest using 1) the Non-Access Stratum (NAS) to send authenticated messages to the device, or 2) Over-The-Air (OTA) UE SIM card provisioning. Because NAS provides message integrity between the eNodeB and UE (mediated by pre-shared keys in the UE SIM card), messages received in this way cannot be spoofed by a (physically) nearby adversary. However, sending alerts over this channel would limit their reception *only* to devices that had established a NAS session. Instead, we recommend using this authenticated channel to send and update a public key that a device should trust. This key should correspond to the private key held by a network operator's Cell Broadcast Center (CBC), which is authorized to broadcast such alerts. Alternatively, the public key distribution can be done using OTA management,⁴

Figure 14. The network provides a machine learning (ML)-based model which characterizes legitimate eNodeBs, and therefore UE can determine whether the alert could have come from a trusted eNodeB or not.

Figure 15. Secure CMAS delivery is guaranteed by adding a signature to alerts. As of May 2019, the FCC mandated to support alert messages up to 360 characters; adding a 64-byte digital signature now becomes applicable for the existing and future wireless emergency alert systems.

which is a well-established technique for updating data on the Universal Integrated Circuit Card (UICC).

To verify this scheme's feasibility, we first stored a public key in a SIM card, assuming that a network operator will provision it. Then we implement the ed25519 digital signature for the Presidential Alert⁷ to sign a 4-byte time stamp along with the CMAS alert message (68 bytes overhead in total). Once a signed message is received, the alert message can be displayed after verifying its signature, as shown in Figure 15. As a result, the UE is not affected by the spoofing attack because it only accepts signed messages.

6. CONCLUSION

In this paper, we have identified the WEA security vulnerabilities over commercial LTE networks and found that a spoofing attack with fake alerts can be made very easily. Specifically, we presented our threat analysis on the spoofing attack and implemented an effective attack system using COTS SDR hardware and open-source LTE software. Our extensive experimentation confirmed that the CMAS spoofing attack could succeed in all tested smartphones in the top four cellular carriers in the U.S. Further, we have proposed potential defenses, from which we believe that completely fixing this problem will require a large collaborative effort between carriers, government stakeholders, and cellphone manufacturers.

References

- 3GPP TR 33.969. Technical Specification Group Services and System Aspects; Study on security aspects of public warning system (PWS) (Release 15), 2018. <http://www.3gpp.org/DynaReport/33969.htm>.
- 3GPP TS 23.041. Technical Specification Group Core Network and Terminals; Technical realization of Cell Broadcast Service (CBS) (Release 15), 2018. <http://www.3gpp.org/dynareport/23041.htm>.
- 3GPP TS 29.168. Technical Specification Group Core Network and Terminals; Cell Broadcast Centre interfaces with the evolved packet core (Release 15), 2018. <http://www.3gpp.org/dynareport/29168.htm>.
- 3GPP TS 31.115. Technical Specification Group Core Network and Terminals; Secured packet structure for (Universal) subscriber identity module (U)SIM toolkit applications (Release 15), 2019. <http://www.3gpp.org/dynareport/31115.htm>.
- 3GPP TS 36.211. Technical Specification Group Radio Access Network; Physical channels and modulation (Release 15), 2018. <http://www.3gpp.org/dynareport/36211.htm>.
- 3GPP TS 36.331. Technical Specification Group Radio Access Network; Evolved universal terrestrial radio access (E-UTRA); radio resource control (RRC) (Release 15), 2018. <http://www.3gpp.org/dynareport/36331.htm>.
- Bernstein, D.J., Duif, N., Lange, T., Schwabe, P., Yang, B.-Y. High-speed high-security signatures. *J. Cryptographic Eng* 2, 2 (2012), 77–89.
- Bui, N., Widmer, J. OWL: a reliable online watcher for LTE control channel measurements. In *ACM All*

Things Cellular (MobiCom Workshop) (November 2016).

9. CellMapper. Cellular coverage and tower map, 2018. <https://www.cellmapper.net>.
10. Chen, X., Jindal, A., Ding, N., Hu, Y.C., Gupta, M., Vannithamby, R. Smartphone background activities in the wild: origin, energy drain, and optimization. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking* (2015), MobiCom'15, Paris, France.
11. Goldsmith, A. *Wireless Communications*. Cambridge University Press, Cambridge, England, August 2005.
12. Gomez-Miguel, I., Garcia-Saavedra, A., Sutton, P.D., Serrano, P., Cano, C., Leith, D.J. srsLTE: an open-source platform for LTE evolution and experimentation. In *ACM WiTECH (MobiCom, Workshop)* (October 2016).
13. Huang, J., Qian, F., Gerber, A., Mao, Z.M., Sen, S., Spatscheck, O. A close examination of performance and power characteristics of 4G LTE networks. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services* (2012), MobiSys'12, Low Wood Bay, Lake District, UK.
14. Li, Y., Peng, C., Yuan, Z., Li, J., Deng, H., Wang, T. Mobileinsight: extracting and analyzing cellular network information on smartphones. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), MobiCom'16, New York City, New York, USA.
15. Lichtman, M., Jover, R.P., Labib, M., Rao, R., Marojevic, V., Reed, J.H. LTE/LTE-A jamming, spoofing, and sniffing: threat assessment and mitigation. *IEEE Commun. Mag.* 54, 4 (April 2016), 54–61.
16. National Public Radio. Officials assess response to camp fire in northern California, 2018. <https://goo.gl/IF12Vo>.
17. NextEPC Inc. Open source implementation of LTE EPC, 2019. <https://www.nextepc.com/>.
18. Nsnam. NS-3: a discrete-event network simulator for internet systems, 2018. <https://www.nsnam.org>.
19. The Washington Post. Cellphone users nationwide just received a 'Presidential Alert.' Here's what to know, 2018. <https://goo.gl/KRfDjf>.
20. Wikipedia. Hawaii false missile alert, 2018. <https://goo.gl/oD9ofx>.
21. Yang, D., Xue, G., Fang, X., Tang, J. Crowdsourcing to smartphones: incentive mechanism design for mobile phone sensing. In *The 18th Annual International Conference on Mobile Computing and Networking* (August 2012), MobiCom'12, Istanbul, Turkey.
22. Yang, H., Huang, A., Gao, R., Chang, T., Xie, L. Interference self-coordination: a proposal to enhance reliability of system-level information in OFDM-based mobile networks via PCI planning. *IEEE Trans. Wirel. Commun.* 13, 4 (April 2014), 1874–1887.

Jihoon Lee, Jinsung Lee, Max Hollingsworth, Eric Wustrow, Dirk Grunwald, and Sangtae Ha ([jihoon.lee-1, jinsung.lee, max.hollingsworth, ewust}@colorado.edu, dirk.grunwald, sangtae.ha), University of Colorado Boulder, Colorado, USA.

Gyuhong Lee ([caixy])@mnd.go.kr, Korea Army Academy, Yeongcheon, South Korea.

Youngbin Im ([ybim])@unist.ac.kr, UNIST, South Korea.

Jinsung Lee is the corresponding author.

This work is licensed under a <https://creativecommons.org/licenses/by/4.0/>

Career & Job Center

The #1 Career Destination to Find Computing Jobs.

Connecting you with top industry employers.

Your next job is right at your fingertips. Get started today!

The new ACM Career & Job Center offers job seekers a host of career-enhancing benefits, including:

Access to new and exclusive career resources, articles, job searching tips and tools.

Gain insights and detailed data on the computing industry, including salary, job outlook, 'day in the life' videos, education, and more with our new Career Insights.

Redesigned job search page allows you to view jobs with improved search filtering such as salary, location radius searching and more without ever having to leave the search results.

Receive the latest jobs delivered straight to your inbox with **new exclusive Job Flash™ emails**.

Get a free resume review from an expert writer listing your strengths, weaknesses, and suggestions to give you the best chance of landing an interview.

Receive an alert every time a job becomes available that matches your personal profile, skills, interests, and preferred location(s).

Visit <https://jobs.acm.org/>