

Aya Fukami University of Amsterdam Netherlands Forensic Institute Netherlands

ABSTRACT

The Replay Protected Memory Block (RPMB) in modern storage systems provides a secure area where data integrity is ensured by authentication. This block is used in digital devices to store pivotal information that must be safeguarded against modification by potential attackers. This paper targets the authentication scheme of the RPMB in three different eMMCs from a major manufacturer. A glitch was injected by sending an electromagnetic pulse to the target chip. RPMB authentication was successfully glitched and the information stored in two target eMMCs was overwritten with arbitrary data, without affecting the integrity of other data.

CCS CONCEPTS

• Security and privacy \rightarrow Security in hardware.

KEYWORDS

RPMB, replay attack protection, glitching, mobile forensics

ACM Reference Format:

Aya Fukami and Richard Buurke. 2024. Keyless Entry: Breaking and Entering eMMC RPMB with EMFI. In *Proceedings of the 17th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '24), May 27– 30, 2024, Seoul, Republic of Korea.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3643833.3656114

1 INTRODUCTION

In modern smart devices, commonly used storage systems often implement a Replay Protected Memory Block (RPMB). The RPMB is a hardware partition in modern storage devices such as Embedded Multi Media Card (eMMC), Universal Flash Storage (UFS), and Nonvolatile Memory express (NVMe).

Writing data to the RPMB requires authentication using a cryptographic hash function, namely a keyed-hash message authentication code (HMAC) using SHA256. The HMAC is calculated over the data frame, excluding some fields such as the HMAC itself. The key used for the calculation is programmed into the storage device only once. In modern embedded devices, because RPMB authentication relies solely on the confidentiality of the pre-shared key, a secure component such as a Trusted Execution Environment (TEE) takes ownership of the RPMB [10]. Figure 1 shows the basic concept of accessing the RPMB.

By issuing a RPMB read request command, anyone can read the content of the RPMB data. Therefore the RPMB is not suitable for



This work is licensed under a Creative Commons Attribution International 4.0 License.

WiSec '24, May 27–30, 2024, Seoul, Republic of Korea © 2024 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0582-3/24/05. https://doi.org/10.1145/3643833.3656114

Richard Buurke Netherlands Forensic Institute Netherlands



Figure 1: RPMB write sequence block diagram

storing confidential information. Rather, the RPMB is commonly used to store information that is immutable for a normal user, for example:

- An anti-rollback counter / version information [14, 16]
- Cryptographic public keys [13]
- Bootloader lock state [31]

As an example, an anti-rollback counter stored in the RPMB can be used as part of the key derivation process. When the device is wiped, because of a factory reset or by exceeding the maximum number of allowed password attempts, the anti-rollback counter in the RPMB is incremented. In this scenario, even when restoring the entire contents of flash user data, key derivation will fail and the encrypted user data cannot be decrypted [14].

Since authentication is performed by a cryptographic hash function using a pre-shared key, integrity of the information stored in the RPMB is guaranteed as long as the secrecy of the key is maintained. This research tries to break this authentication scheme, enabling an attacker to overwrite data in the RPMB without knowledge of the pre-shared key. We targeted the RPMB in eMMCs due to their availability, and usage in a wide variety of embedded products, such as smartphones [14], IoT devices [16] and automotive systems [24].

Several techniques for attacking RPMB authentication exist, and will be covered more extensively in Section 3. For this experiment fault injection (FI) was applied, which is an umbrella term for a collection of techniques that aims to introduce faults, or glitches, into a device leading to unintended behavior. This can be achieved through software [21, 30] or hardware methods [26, 28, 29]. Commonly used hardware methods are:

- Shorting the power supply (crowbar glitching, or voltage fault injection (VFI))
- Introducing electromagnetic pulses (EMFI)
- Illumination with a laser beam (LFI)
- Changing the clock signal

Fault injection can be used to circumvent security checks in software whilst the program itself does not contain any (known) vulnerabilities. Unintended behavior that can aid the attacker is also referenced as a fault primitive, such as skipping instructions or corrupting CPU register values.

By applying EMFI on an eMMC device we were able to write arbitrary data to the RPMB, without knowledge of the pre-shared key and thereby breaking the RPMB authentication scheme which compromises the integrity of all stored data. We then applied the same method to a different model eMMC device which was also successfully compromised.

This research makes the following contributions:

- We show how EMFI can be applied on an eMMC device to enable writing of arbitrary data to the RPMB, using commercially available tools.
- A detailed description on how to characterize the susceptibility to EMFI, on eMMC controllers from a single manufacturer, through arbitrary code execution is provided.
- Advantages and disadvantages of applying EMFI, compared to other methods for breaking RPMB authentication are outlined.

The structure of the rest of this paper is as follows. In Section 2, we discuss the threat model related to our research. Section 3 discusses related work to this research. The targets and physical FI setup are described in Section 4. Section 5 outlines how the firmware of our target devices was obtained, and reversed engineered. This section also describes how this information was used to run our own fault observer code on the target, to characterize the susceptibility to EMFI. The process for successfully applying EMFI on two of the targets, is detailed in Section 6. The impact of this attack is discussed in section 7 before concluding in section 9.

2 THREAT MODEL

Because the RPMB is a type of authenticated storage, it is often used for storing non-confidential information that requires integrity protection. Examples include version information used in anti-rollback mechanisms [9, 14, 31], bootloader lock state [9, 31] and asymmetric public keys [9, 13]. Modification of this type of information by an unauthorized attacker can result in an increased attack surface, or in some cases, a fully compromised system. Fukami, et al. (2024) [14] showed that it is possible to decrement an anti-rollback counter stored in the RPMB to return a wiped Android device to a non-wiped state. Anti-rollback mechanisms are also often used to prevent software components from being downgraded to a vulnerable state [2, 3].

Modern Android systems are commonly based on the ARM Trusted Firmware-A reference design, which implements a chainof-trust (CoT) for the boot process [5]. Figure 2 shows a simplified representation of a typical secure boot implementation. Devices based on the ARM Cortex-A architecture feature a system-on-chip (SoC) which can switch between multiple contexts. The normal operating system (OS) runs in the so-called "normal world" while a dedicated second operating system is executed in the "secure world". This separation provides an additional layer of security since the OS running in normal world cannot access memory and resources used by the secure world.

After powering on the device, the first boot stage, embedded in the ROM of the SoC, is executed. The secondary boot stage is loaded from flash and verified using a public key stored in eFuses [23]. Each subsequent boot image is verified before being executed. The Android Bootloader (ABL) implements Android Verified Boot (AVB), and is responsible for starting the Android operating system. If the bootloader is locked, a public key verifies the integrity of the boot image, and hashtrees for the system and vendor partitions [1]. Some vendors, such as NXP [9], recommend storing the bootloader state and AVB key in the RPMB area of an eMMC. By manipulating this information an attacker can unlock the bootloader or modify and resign the Android operating system, resulting in code execution with EL1 privileges in the normal world.

3 RELATED WORK

Recovering the secret key has been the main focus of security research related to HMAC authentication. Jeong et al. showed that they could successfully recover the key by injecting faults and reducing the number of computational rounds to HMAC [19]. Belenky et al. published their research on a side channel attack against HMAC. The authors successfully simulated recovery of



Figure 2: ARM Trusted Firmware-A secure boot overview

Target number	1	2	3
eMMC version	4.41	5.1	5.0
Part no.	0xc23 (Cortex-M3)	0xc27 (Cortex-M7)	0xc23 (Cortex-M3)
Architecture	0x0F (ARMV7-M)	0x0F (ARMV7-M)	0x0F (ARMV7-M)
Variant	0x02	0x01	0x02
Revision	0x00	0x01	0x00
MPU enabled	No	No	No
VTOR	0x40000	0x60000000	0x40000

Table 1: Target chip properties

the key through a template attack [7]. Our research focuses on the application of HMAC in a real-world device, where the HMAC authentication itself needs to be skipped.

Western Digital published a white paper on vulnerabilities in the eMMC RPMB [34], and suggested that by performing a "manin-the-middle" attack, one can trick the host system to believe that writing the data was successful, whilst in the background different data is being written. Their attack is effective if targeting a "live" device, where an attacker can monitor the command issued by the host. In our scenario, arbitrary data is written to the RPMB without relying on the behavior of the host device.

Fault injection attacks against mobile or IoT devices have been widely performed by multiple researchers, as reported by Gangolli et al. and Shepherd et al. [15, 29]. In their research, SoCs are the main target to introduce glitches in order to skip the cryptographic authentication and gain root privileges on the target device.

Multiple publications discuss attacking ARM Cortex-M microcontrollers. Bozzato et al. showed that it is possible to improve the effectiveness of voltage glitching attacks by injecting an arbitrary waveform instead of shorting the power supply [8]. They validated their approach on multiple microcontrollers including an STM32F103 (Cortex-M3) and STM32F373 (Cortex-M4). Ruminot et al. focused on improving voltage glitching attacks by reducing the voltage supplied to the target just outside of normal operating parameters [27]. Werner et al. created a tool that is able to find fault injection parameters in an automated manner, by combining characterization data with a simulated model [33]. The authors applied their approach to a Cortex-M4 microcontroller, using a dual-laser fault injection setup.

Although the aforementioned research targets similar devices as in our research, they do not use EMFI. Instead, they opt for VFI or LFI. Furthermore, research by Ruminot et al. [27] and Werner et al. [33] did not circumvent any real-life security mechanisms.

To the best of our knowledge, this is the first work to perform EMFI against eMMC RPMB authentication. The hardware designs of the targets are proprietary and unknown. We also did not have access to the source code, or reference implementations, of the firmware.

4 EXPERIMENT SETUP

4.1 Electromagnetic Fault Injection

We chose to perform EMFI for our attack scenario since it does not require any additional hardware modifications and provides localized injection of faults. LFI requires thinning of the package to expose the internal transistors, we therefore did not consider this method. The external clock signal might be susceptible to clock glitching, however we assumed that it was not directly connected to the application processor (AP) embedded in the controller. Since an eMMC device is powered from an external source, the core voltage of the controller (Vddi) and memory peripherals (Vcc) can be trivially manipulated. Therefore it might also be possible to apply voltage glitching to these devices.

4.2 Target Selection

The goal of this research is to bypass the cryptographic authentication of the RPMB in an eMMC. Prior to the actual attack, we decided to run a profiling process in order to identify the location of the chip where it is most susceptible to the EMFI attack. For this purpose, we selected eMMCs from a single manufacturer where we have access to their firmware. We selected 3 eMMC devices shown in Table 1 on availability in our forensic lab. Upon request from the manufacturer, we refer to these devices as Target 1, 2, and 3, respectively. Prior research suggests that our target devices contain a Cortex-M micro controller, and that the firmware of these devices can be read using vendor-specific commands [6]. In 2018, Avraham demonstrated that it was possible to read/write memory regions of an eMMC controller, embedded in a mobile phone, using proprietary vendor commands [6]. The author subsequently released the proof-of-concept code on his public repository [25]. We observed that these vendor commands are still implemented in an eMMC that is currently in mass-production. Therefore we assume that they are still applicable to a wide range of devices.

The firmware of each target device was extracted by using an EasyJTAG flasher box [35], which supports the required vendor-specific commands.

4.3 Glitching Setup

Each target eMMC was mounted on a custom breakout adapter in order to avoid the need for reballing and resoldering. The socket for the adapter is soldered onto a custom-made printed circuit board (PCB), which is mounted onto the XYZ table (Genmitsu 3018 PROver V2) with a custom-made fixture for precise positioning and repeatability. Electromagnetic pulses were injected by a NewAE ChipSHOUTER (CW520) [17], using a one millimeter clockwise winded probe. Figure 3 shows an overview of the setup. Communication with the target eMMC was implemented using a programmable I/O (PIO) based state machine, running on a Raspberry Pi Pico. The ChipSHOUTER is also triggered by the state machine program, using a fixed delay after sending the eMMC command. The rest of the fault injection process was orchestrated by software running on a Raspberry Pi 4.

A Tektronix DPO7354C oscilloscope was used to monitor eMMC communication and measure EM emissions from the application processor (AP). EM measurements were used to determine the timing of eMMC operations, as described in Section 6.1. However any four channel scope with around 200MHz bandwidth should suffice, such as a PicoScope 3000 Series.

Whilst industry level FI equipment is available, cost effective off-the-shelf products were used as much as possible for this attack setup.



Figure 3: Target chip mounted to a custom PCB, using an adapter, while applying an EM pulse using a ChipSHOUTER (CW520)

5 CHARACTERIZATION

Given that EMFI enables us to localize the attack, we first needed to identify the most vulnerable location of the target eMMC. We therefore acquired and analyzed the firmware of each target device, in order to implement a simple fault observer program that could run on the target device. The fault observer enabled us to monitor the effect of the EM pulse at different locations.

5.1 Device Identification

We extended the *mmc-utils* program [22] to include the vendorspecific commands discussed in Section 4.2. We assumed that the flash controller was based on the ARM architecture, and so our version of *mmc-utils* was used to read the System Control Block (SCB), at offset 0xE000ED00, and the Memory Protection Unit (MPU) configuration, at offset 0xE000ED90. Table 1 combines information from the SCB, MPU configuration and Device identification (CID) register.

First, we noticed that the MPU is disabled for all chips. According to the *ARMv7-M Architecture Reference Manual* [4], the default system memory map is used when the MPU is disabled. By default, the *Code, SRAM* and *RAM* memory segments are mapped readable, writeable and executable. The vector table offset register (VTOR) points to the main vector table of the device. The second entry in this table holds the address of the reset handler, which is the entry point for our ROM code.

Arbitrary Code Execution. The code section also holds the vector table for standard eMMC commands and can therefore be overwritten using vendor-specific commands. For Target 2 this table was located in the RAM segment, but the same principles apply. To gain arbitrary code execution, the payload was first written to an unused memory region, and the entry for CMD8 (SEND_EXT_CSD) in the vector table was updated with the address of our routine.

<pre>void fault_observer(void) { uint32_t total_iterations; uint32_t value; uint32_t j; uint32_t i; extcsd *ext_csd;</pre>
ext_csd = PTR_EXT_CSD;
<pre>total_iterations = 0;</pre>
value = 0;
j = 0;
do {
j = j + 1;
i = 0;
do {
value = value + 7;
i = i + 1;
} while ((int)i < 62500);
<pre>total_iterations = total_iterations + i;</pre>
<pre>} while ((int)j < 4);</pre>
<pre>ext_csd->total_iterations = total_iterations;</pre>
ext_csd->value = value;
(*(code *)CMD8)();
return;
}

Listing 1: Fault observer implementation

Our fault observer implementation was directly written in assembly. Listing 1 shows the decompilation of this code for readability. It consists of a nested for loop that increments an unsigned integer value for every iteration. The total number of iterations and incremented value are written to the beginning of the extended CSD register, which according to the current JEDEC standard [18] is unused. Finally the original CMD8 routine is executed, returning the contents of the extended CSD register. By checking the stored values, it is possible to determine if the controller was affected by the EM pulse. This approach worked for all targets.

5.2 Profiling

By introducing EM pulses at different positions of the chip while the fault observer code was running on the target, we determined which areas of the device were the most affected. Based on the return values of the fault observer, each attempt was categorized as *Normal, Crash*, or *Glitch*. If the fault observer returned the expected value, the result was categorized as *Normal*. If the response from the target was all 0x00 or 0xFF, then the result was categorized as *Crash*, since at this point the target chip cannot behave normally unless a hard reset is performed. If the result from the fault observer

18

19

24



Figure 4: Profiling results for each chip using the fault observer. The lighter the color, the more susceptible the chip is for the categorized result at that location.

was different from the expected value, but the target chip still operated normally, then it was categorized as *Glitch*. The probe was positioned using a 1mm by 1mm grid, overlayed on the target chip. A glitch attempt was performed at each position for 25 iterations in total. An EM pulse with a strength of 200V and length of 100ns, using a clockwise winded probe with a 1mm diameter, was used for profiling. A heatmap was created of each categorized result (Figure 4).

As shown in Figure 4, it is clear that Target 1 is susceptible to the glitching attack at multiple locations. Whereas Target 2 and 3 can be only glitched or crashed if the EM pulse is sent around the perimeter of the chip. Nevertheless, through profiling, it is clear that a fault can be injected using EM pulses, to affect the code running on the eMMC controller for each target. The heatmap was then compared with the internal structure of each target, using X-rays, and the location of the flash controller in each target chip was identified. Figure 5 shows the X-ray image of each target. Based on the bonding wires of the internal controller, the estimated location of the controller is highlighted with an arbitrary color.

For Target 1, the location of the "hot spot", where the fault observer can be successfully glitched, matches the location of the controller (Figure 4a and 5a). Therefore it is possible to assume that



(a) Target 1 (b) Target 2 (c) Target 3 Figure 5: X-ray inspection of target chips, highlighting the physical location of the controller

Tał	ole	2:	Gl	itcl	hing	pa	ram	eters	for	each	ı tar	get
-----	-----	----	----	------	------	----	-----	-------	-----	------	-------	-----

Target	1	2	3
X position (mm)	154	159.8	159.8
Y position (mm)	-62.5	-58.1	-58.5
Z position (mm)	-25.7	-25.7	-25.7
Pulse Voltage (V)	200	200	200
Pulse Duration (ns)	100	100	100

the fault is directly injected to the logic of the controller. On the other hand, the hot spot of Target 2 does not match the location of the controller (Figure 4b and 5b). Rather, crashes or glitches appear to happen when the EM pulse is sent directly on top of one of the bonding wires. This bonding wire is, to our best knowledge, related to VCC or GND. It is also worth mentioning that the controller of Target 2 is located under the flash memory dies. Therefore the distance between the EM probe and the controller of this target is larger than Target 1 or Target 3. The fault observing procedure was repeated with EM pulses at a higher voltage, however the result were the same as shown in Figure 4b. The hot spot of Target 3 matches the location of the controller, however the glitching rate is much lower (less than 10 %), compared to Target 1 (around 30 %), for the best location.

Subsequently, the optimal glitching parameters for each target chip were determined. Different voltages and lengths of the EM pulse were tried at the most susceptible location of each chip (Location x=6, y=4 for Target 1, Location x=10, y=0 for Target 2, and Location x=10, y=1 for Target 3, as shown in Figure 4). The pulse voltages and lengths were selected between 150V and 250V, and between 40ns and 1000ns, respectively. The parameters were randomly selected while repeating the operation for 1500 times. We 10 observed that Target 2 and 3 are more susceptible to crashing if the voltage exceeded 200V. However, the length of the pulse did not 12 seem to affect the result. The results were uniformly distributed ¹³ regardless of the length of the pulse. Therefore, we chose the glitching parameters as shown in Table 2. The X, Y, Z position values are based on our XYZ table setup.

5.3 Firmware Reverse Engineering

In order to gain better insight into the RPMB authentication implementation, we reverse engineered the firmware from Target 1. All available memory areas were dumped from the chip, including the boot ROM and main ROM code, using the aforementioned vendor commands.

The address of the reset handler was the start of our reverse engineering effort, then the location of the command handler loop was determined. The firmware uses a vector table located in SRAM, or RAM segment in case of Target 2, for all standard eMMC commands.

Fault Modeling. Commands CMD24 (WRITE BLOCK) and CMD25 (WRITE_MULTIPLE_BLOCK) are handled by the same function and implements all RPMB functionality. This function was further analyzed in order to understand how RPMB key authentication could be circumvented using fault injection.

Listing 2 is the decompiler output for the routine that checks the HMAC of an RPMB write request.

```
uint32_t rpmb_check_hmac(void *hmac,uint32_t
\rightarrow length) {
  uint32_t i = 0;
  if (length + 3 >> 2 != 0) {
    do {
      if (*(int *)((int)hmac + i * 4) != *(int
       \rightarrow *)(CORRECT_HMAC + i * 4 + 0x60)) {
        return 0;
      }
      i = i + 1;
    } while (i < length + 3 >> 2);
  }
  return 1;
}
```



8

11

The routine checks the HMAC from the eMMC packet against a pre-calculated HMAC stored in a hardware register, four bytes at a time. It returns 1 if the HMAC is valid, otherwise it returns 0. This routine was also implemented in Target 3. The following fault injection possibilities were observed:

- (1) If the length argument is set to 0 the check is skipped in its entirety
- (2) If the register r0 is set to any non-zero value, the ROM assumes the HMAC was valid
- (3) If we skip the call to *rpmb_check_hmac* entirely, the verification will succeed, because r0 contains a pointer to the provided HMAC (non-zero value)

6 GLITCHING RPMB AUTHENTICATION

6.1 Glitching Setup

Based on the results presented in Section 5.2, we hypothesized that it should be possible to skip the RPMB HMAC authentication routine using EMFI. When writing data to the RPMB, a JEDECdefined data packet needs to be sent to the target chip. The data packet is 512-bytes long, and should contain the data to be written, Nonce, the write counter, address, block count, request message, and an HMAC calculated over this data [18]. Upon receiving the data packet, the controller in the eMMC calculates the HMAC using the previously programmed key. If the HMAC matches the received one, the controller starts the data writing operation. The command sequence of the RPMB write routine is shown in Figure 6a. Commands and data in white boxes are sent from the host to the target. Orange boxes indicate responses from the target. RPMB authentication is most likely performed during the time indicated by the red box, thus the timing for performing the fault injecting attack is critical.

In order to precisely identify the attack timing window, the EM radiation emitted from the controller was measured at the moment when the RPMB data packet was sent to the target. Figure 6b shows the captured waveform. The captured timing matches the timing window indicated by the red box in Figure 6a. DAT0 line is shown in magenta, CLK line in green, and EM emission in blue. As shown in the figure, during and after the data is sent to the target, the controller keeps performing internal operations. The CRC status of the sent data (positive = "010") is sent on the DAT0 line from the controller, synchronized with the CLK signal, as defined by JEDEC Standard [18]. This operation is followed by DAT0 line pulled low, which signals that the target device is busy performing the internal computations. Even after the DAT0 goes back to high, the operation continues on the controller. Since no other commands are issued to the target, we made an assumption that the HMAC verification is most likely performed during this timeframe, and that the observed EM emission comes from the controller processing the RPMB data and performing the HMAC computation.

After sending the RPMB result request, the host can request reading the result register value, which is initiated by a read command (CMD23 (SET_BLOCK_COUNT) and CMD18(READ_MULTIPLE_ BLOCK), shown in Figure 6a). Table 3 shows the list of the result register values defined by JEDEC [18]. Whilst more values are defined, only the relevant values in this setup are listed in Table 3.



(b) Waveforms when data with wrong HMAC were sent to the target. The timing matches the timing window shown in red in Figure 6a

Figure 6: RPMB write command scheme and captured EM emissions

Table 3: Partial list of RPMB operation result register values

Value	Results					
0x00	Operation OK					
0x01	General failure (Multiple errors have occurred)					
0x02	Authentication failure (HMAC mismatch)					
0x03	Counter failure					
0x04	Address failure					

The returned register values were monitored to determine if the fault was successfully injected into the RPMB authentication procedure. If the returned value was 0x02, the target was determined to be responding normally. Because we intentionally sent an incorrect HMAC value, this is the expected result. If the target responded with value 0x00, the attack had successfully skipped the authentication procedure. If any other value was returned, then it was determined that an error occurred during the RPMB authentication procedure. If all of the responses were 0x00 or 0xff, the target was presumed to have crashed, since a hard reset is required to bring the chip back to a normal working state.

The setup described in Section 4.3 was reused, however the Raspberry Pi Pico was reprogrammed to communicate with the RPMB on the target eMMC. The JEDEC eMMC Electrical Standard [18] was followed to send the required command sequence. First, the RPMB was programmed with an arbitrary key. Then the first block of the RPMB (256 bytes) was programmed with random values. Subsequently the RPMB write routine was repeated with the wrong HMAC value and the EM pulse was introduced during the above mentioned timing window. After completing the write operation and waiting for the busy signal to be cleared, the RPMB result



Figure 7: Returned result register value after RPMB authentication glitching

request was sent to the target, followed by the result reading command. The 200V EM pulse with a length of 100ns was injected at 10ns granularity during the target timing window. The trigger signal was generated when the last bit of the data packet was sent. The time-frame between the trigger and the end of the busy signal is around 119us for Target 1, and 113us for Target 3. The procedure was repeated until the target responded with the "Operation OK" register value, and the correct response. Additionally, in case the system data got corrupted, vendor-specific commands were implemented in the PIO state machine of the Raspberry Pi Pico, which enabled us to restore the system data.

6.2 Results

The glitching campaign targeting RPMB authentication was executed exclusively on Target 1 and 3. During the profiling campaign, Target 2 became non-responsive, resulting in the unavailability of samples for analysis. Figure 7a shows the returned result register values over time on Target 1. The x-axis shows the timing where the EM pulse was injected, and the y-axis shows the actual returned value of the result register. The timing is shown as a delay since the last bit of write data was sent to the target. Responses with 0xFFFF and 0x0000 occur mostly when the target has crashed. Register values 0x01, 0x03 and 0x04 are returned multiple times by introducing the EM pulse. The assumption is that the internal operations were corrupted through EMFI, while the rest of the RPMB routine continued executing on the controller. Additionally, the write counter value was also overwritten with an unexpected value several times on Target 1, because it is stored in SRAM. We had to restore this value by using vendor-specific commands before continuing the glitching campaign. Nevertheless, the RPMB authentication was successfully bypassed at the timing shown as red dots in Figure 7a.

After successfully glitching Target 1, we repeated the same procedure on Target 3 by using the parameters defined in Table 2. Figure 7b shows the returned result register value. Similar to what was observed on Target 1, 0x01, or "General failure" was often observed at the early stage of the RPMB authentication. At the same time, Target 3 crashed more often than Target 1, where the response packet was filled with 0xFFFF. Nevertheless, the RPMB authentication was successfully bypassed when the EM pulse was injected around 112us, as shown in Figure 7b. Unlike Target 1, no data corruption was observed during the glitching campaign on Target 3.

The critical timing moment is from 117.72us to 118.30us for Target 1, and from 112.3us to 112.50us for Target 3 since the trigger (after the last bit of the data packet was sent from the PIO), where allegedly one of the scenarios suggested in Section 5.3 was executed. On both targets, those timings are around the end of the busy signal indicated on the DAT0 line. After successful glitching, the counter value was incremented by one on both targets, which means that the controller acted as if the correct HMAC was received and proceeded with writing the received value. Post reading of the RPMB value confirmed that the sent data was successfully stored. Therefore we concluded that the RPMB authentication can be successfully skipped through EMFI if it is performed at the correct timing.

6.3 Integrity of Non-Volatile Data

Since the critical timing to skip the RPMB authentication check is identified, we repeated the experiment with a new chip. Prior to the experiment, arbitrary data is written into the user data area. Before starting the glitching campaign, all physical data was dumped and hashed using SHA-256. Figure 8 shows the mounting of an eMMC on a custom-made eMMC-SD adapter. The adapter was connected to a PC using an SD-card slot. The data is imaged using the dd command on a Linux operating system (OS).



Figure 8: Target device mounted on an SD-eMMC Adaptor

After completing the non-volatile data extraction, the target chip was mounted on the glitching setup, and the EM pulse was injected only at the critical timing of the RPMB authentication, identified above. Both on Target 1 and 3, bypassing the RPMB authentication was successful in less than 10 tries. After successfully bypassing the RPMB authentication, the user data area was again extracted. The SHA-256 hash value of the extracted data successfully matched the value computed before the glitching attack. Additionally, the RPMB data was only overwritten at the specified block address, keeping the remaining block data unchanged. Therefore we conclude that EMFI attacks against RPMB authentication can be conducted, while preserving integrity of the stored data, by introducing pulses using predetermined parameters, including the timing.

It must be noted, however, that repeatedly sending EM pulses to the same device does seem to increase the chance of data corruption in the user data area of the eMMC. We repeated the glitching campaign on Target 1 and Target 3 for additional 100 consecutive times using our setup. We observed multiple data corruptions in the user data area of both targets, while the RPMB area did not seem to be affected. Any corruptions could easily be reverted using the write command to the affected data sectors. However this observation emphasizes the importance of creating a data backup beforehand, and checking the integrity of the data after a successful glitching attempt.

7 DISCUSSION

7.1 Arbitrary Code Execution

It is possible to run our own fault observer code on the target during the characterization phase due to vendor-specific commands provided by the manufacturer. Although convenient, this is not a requirement for device characterization. The response from the eMMC controller, after requesting the result for an RPMB write request, includes a result register [18], as noted in Table 3 in Section 6.1. This register indicates which operation failed while processing the request (e.g. HMAC mismatch, expired write counter, etc.). It was observed that applying EMFI to the target while it processes an RPMB write request, changes this value. Therefore other methods can be used for characterization, that do not rely on arbitrary code execution.

7.2 Timing

Our initial setup relied on the Linux kernel eMMC driver for communication with the target device. However, the scheduling properties of a non real-time operating system, makes it unsuitable for FI purposes. Using the Linux kernel driver also lacks the required fine-grained control needed for communication. The kernel driver periodically sends commands, that are not controlled by the user, to check the status of the device. Therefore programmable input/output (PIO) on the Raspberry Pi Pico was used instead. A benefit of this approach is that the timing can be precisely controlled. However, this also means that a deep understanding of the eMMC protocol is required to correctly implement the PIO state machines, for communication with the target device.

Although various open source implementations are available [12, 36], they eventually did not meet the requirements for this experiment, and therefore a custom implementation was used.

7.3 Data Corruption

By sending high-voltage EM pulses to the target device, we observed both user data and system data corruption. Especially through profiling, where EM pulses are shot directly on the flash memory dies, some sectors became unresponsive. In some cases, the target eMMC became completely unresponsive, preventing further analysis. These results show that repeating EM pulses are highly destructive on eMMCs. Therefore it is important to perform the profiling using a reference device to identify the correct timing and location. Furthermore, creating a backup of the stored data before executing the EMFI attack is highly recommended.

7.4 Real World Applications

Even though use of the RPMB, to the best of our knowledge, has only seen limited applications in smartphones (e.g. storing an antirollback counter [14]), use of this type of storage seems to be embraced by the automotive industry. For example, U-Boot is a popular bootloader used in embedded devices, and stores anti-rollback counters and the bootloader lock state in the RPMB [31], when Android Verified Boot (AVB) is configured to use OP-TEE.

According to NXP, their I.MX family of processors is used by most large car manufacturers [11]. NXP also supports Android Automotive, that uses Trusty as the operating system of choice to run in the TEE [24]. According to the *i.MX Android Security User's Guide* from NXP [9], the RPMB is used to store anti-rollback counters, bootloader lock state and AVB public key, which is used for verifying integrity of system images. The RPMB key itself is encrypted and decrypted in the TEE by Trusty. The Digi ConnectCore 8X system-on-module (SOM), which is designed around the NXP I.MX 8X processor, and uses eMMC storage, also stores the AVB public key in the RPMB [13].

It therefore seems that the RPMB is used as a critical component in the secure boot implementation of a wide variety of automotive products. Compromising the RPMB means that an attacker will be able to rollback potentially vulnerable versions of software, unlock the bootloader or re-sign system images. Resulting in the attacker gaining root privileges in the Android operating system.

Whilst an eMMC chip does provide tamper resistant storage in the form of an RPMB, any mitigations against fault injection have not been encountered during our experiments. Furthermore, the total cost of re-creating our FI setup is less than USD\$7.000, meaning that applying EMFI is in reach for most attackers.

7.5 Mitigations

As discussed in Section 5, publicly known vendor-specific commands were used to achieve code execution on all devices, and run our fault observer routine. In the case of Target 1 and 3, this already breaks the security of the RPMB, since the firmware can be patched. The firmware of Target 2 was not fully reversed.

Applying EMFI on an eMMC chip requires physical access to the device. Depending on the threat model being used, this might be a valid concern. One way to increase the difficulty of successful FI attacks is to implement mitigations in software as described by van Woudenberg and O'Flynn [32] (e.g. double checking critical data, using non-trivial constants). As mentioned in Section 5.3, the implemented HMAC validation routine shown in Listing 2 only fails when returning 0. This requirement is trivial to achieve since the CPU register holding the return value (r0) holds a pointer to the HMAC, and thus is non-zero before the function is called. Requiring a return value with a large Hamming distance (i.e. 0xA5C3B4D2) significantly increases the attack complexity. A large number of bit flips is needed to end up with the correct return value.

Validation of the HMAC is critical, if circumvented the integrity of the RPMB is fully compromised. Therefore the HMAC should be checked multiple times. Preferably a random delay should be added between both checks. This requires the attacker to insert multiple glitches with a non-constant delay in between.

The implemented routine is also vulnerable to a timing attack. The function returns as soon as the comparison fails. A possible mitigation would be to always check the entire HMAC, ensuring that the comparison uses constant time.

Hardware mitigations against fault injections attack could also be considered. For example, Jiang et al. showed, using a simulated model, how machine learning can be applied to detect voltage glitching attacks in low power circuits [20]. Ruminot et al. proposed adding a potentiometer, with a random resistance whenever the device is started, coupled with a capacitor, in order to mitigate the effect of voltage glitching attacks [27]. The authors suggest integrating this circuit within the same package as the microcontroller.

7.6 Future Work

This research focused on attacking the RPMB embedded in eMMC devices. However succeeding technologies, such as UFS and NVMe, also include an RPMB implementation. UFS is the storage solution of choice for the current generation of smartphones making it an interesting subject for future research.

As described in Section 2, by compromising the integrity of the RPMB data, an attacker could take over the target system. Expanding this research into a consumer device, such as automotive devices, represents a promising avenue for further exploration.

8 RESPONSIBLE DISCLOSURE AND RESOURCES

We have reported our findings to the manufacturer of the target devices, and our report was received by the manufacturer on 11 February, 2024. They thoroughly investigated the issue and identified multiple vulnerabilities. At the time of this publication, the manufacturer is developing a patch for this matter, and we are obligated not to disclose any manufacturer-related information until 25 September, 2024. Meanwhile, our PIO code for Raspberry Pi Pico, fault injection code, and the custom PCB design file, and other related resources are available at https://github.com/topig/RPMB_Glit ching.

9 CONCLUSION

We have shown that it is possible to circumvent the RPMB authentication scheme and write arbitrary data, by applying EMFI on eMMCs from a major manufacturer. The RPMB functionality of non-volatile storage devices is often considered to store data that requires to be immutable to the user. The potential for bypassing RPMB authentication and manipulating its data raises concerns regarding the integrity of security mechanisms, including anti-rollback protection, bootloader lock state, and signature verification. Widely available, off-the-shelf commercial components were used for these experiments. The total cost of the equipment, needed to re-create the setup, is less than USD\$7.000, which makes it a feasible option for most attackers.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their time and the insightful reviews. The authors would also like to thank Ronald van der Knijff and Zeno Geradts for supervising this research, Nils Wiersma for his technical support and sharing his software, Erik van Dijk for manufacturing the fixture for the XYZ table, Martien de Jongh for providing us the eMMC adapter, and Susan Laraghy for initial proof-reading and revising the article.

WiSec '24, May 27-30, 2024, Seoul, Republic of Korea

REFERENCES

- [1] 2024. Android Verified Boot 2.0. https://android.googlesource.com/platform/ external/avb/+/main/README.md [Online; accessed 1. Apr. 2024].
- [2] 2024. Trust Issues: Exploiting TrustZone TEEs. https://googleprojectzero. blogspot.com/2017/07/trust-issues-exploiting-trustzone-tees.html [Online; accessed 29. Mar. 2024].
- [3] 2024. Verifying Boot. https://source.android.com/docs/security/features/ verifiedboot/verified-boot#rollback-protection [Online; accessed 29. Mar. 2024].
- [4] Arm. 2021. ARMv7-M Architecture Reference Manual. Arm (Dec. 2021). https://developer.arm.com/documentation/ddi0403/latest
- [5] Arm. 2021. Trusted Board Boot Requirements Client (TBBR-CLIENT) Armv8-A. Arm (Dec. 2021). https://developer.arm.com/documentation/den0006/latest
- [6] Oran Avraham. 2018. eMMC hacking, or: how I fixed long-dead Galaxy S3 phones. https://media.ccc.de/v/34c3-8784-emmc_hacking_or_how_i_fixed_ long-dead_galaxy_s3_phones [Online; accessed 23. Jan. 2024].
- [7] Yaacov Belenky, Ira Dushar, Valery Teper, Hennadii Chernyshchyk, Leonid Azriel, and Yury Kreimer. 2021. First Full-Fledged Side Channel Attack on HMAC-SHA-2. 31–52. https://doi.org/10.1007/978-3-030-89915-8_2
- [8] Claudio Bozzato, Riccardo Focardi, Francesco Palmarini, et al. 2019. Shaping the glitch: optimizing voltage fault injection attacks. *IACR transactions on cryptographic hardware and embedded systems* 2019, 2 (2019), 199–224.
- [9] NXP B.V. 2020. i.MX Android[™] Security User's Guide. https: //community.nxp.com/pwmxy87654/attachments/pwmxy87654/imxprocessors/167888/2/i.MX_Android_Security_User's_Guide.pdf [Online; accessed 5. Feb. 2024].
- [10] Liang Cai. 2019. Guard Your Data With the Qualcomm Snapdragon Mobile Platform. https://www.qualcomm.com/content/dam/qcomm-martech/dmassets/documents/guard_your_data_with_the_qualcomm_snapdragon_ mobile platform2.pdf [Online; accessed 4. Feb. 2024].
- [11] Carl Chien. 2017. I.MX in automotive. https://www.nxp.com/docs/en/supportinginformation/BL-Micro-i.MX-in-Automotive-Carl-Chien.pdf [Online; accessed 5. Feb. 2024].
- [12] democloid. 2024. pico-sdio-example. https://github.com/democloid/pico-sdioexample [Online; accessed 4. Feb. 2024].
- [13] Digi International Inc. 2024. Secure boot flow | ConnectCore 8X. https://www.digi.com/resources/documentation/digidocs/embedded/android/ dea11/cc8x/android-trustfence_r_secure-boot-flow [Online; accessed 2. Feb. 2024].
- [14] Aya Fukami, Richard Buurke, and Zeno Geradts. 2024. Exploiting RPMB authentication in a closed source TEE implementation. Cryptology ePrint Archive, Paper 2024/180. https://doi.org/10.1016/j.fsidi.2023.301682 [Online; accessed 9. Feb. 2024].
- [15] Aakash Gangolli, Qusay H. Mahmoud, and Akramul Azim. 2022. A Systematic Review of Fault Injection Attacks on IoT Systems. *Electronics* 11, 13 (2022). https://doi.org/10.3390/electronics11132023
- [16] Dennis Giese and Guevara Noubir. 2021. Amazon echo dot or the reverberating secrets of IoT devices. In Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks. 13–24.
- [17] NewAE Technology Inc. 2020. ChipSHOUTER Kit. https://www.newae.com/ products/nae-cw520 [Online; accessed 6. Feb. 2024].
- [18] JEDEC Solid State Technology Association. 2015. Embedded Multi-Media Card (e·MMC) Electrical Standard (5.1). JEDEC Standard JESD84-B51. https://www. jedec.org/system/files/docs/JESD84-B51.pdf
- [19] Kitae Jeong, Yuseop Lee, Jaechul Sung, and Seokhie Hong. 2013. Security analysis of HMAC/NMAC by using fault injection. *Journal of Applied Mathematics* 2013 (01 2013). https://doi.org/10.1155/2013/101907

- [20] Wei Jiang et al. 2022. Machine Learning Methods to Detect Voltage Glitch Attacks on IoT/IIoT Infrastructures. *Computational Intelligence and Neuroscience* 2022 (2022).
- [21] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. 2014. Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors. ACM SIGARCH Computer Architecture News 42, 3 (2014), 361–372.
- [22] mhei. 2024. mmc-utils. https://github.com/mhei/mmc-utils [Online; accessed 23. Jan. 2024].
- [23] Ryan P Nakamoto. 2016. Secure boot and image authentication. Qualcomm Technologies Inc., San Diego (2016).
- [24] NXP. 2024. Android Automotive for i.MX Applications Processors. https://www.nxp.com/design/design-center/software/embedded-software/imx-software/android-automotive-os-for-i-mx-applications-processors: ANDROID-AUTO [Online; accessed 5. Feb. 2024].
- [25] oranav. 2024. i9300_emmc_toolbox. https://github.com/oranav/i9300_emmc_ toolbox [Online; accessed 23. Jan. 2024].
- [26] Pengfei Qiu, Dongsheng Wang, Yongqiang Lyu, and Gang Qu. 2019. VoltJockey: Breaching TrustZone by Software-Controlled Voltage Manipulation over Multicore Frequencies. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (London, United Kingdom) (CCS '19). Association for Computing Machinery, New York, NY, USA, 195–209. https: //doi.org/10.1145/3319535.3354201
- [27] Nicolás Ruminot, Claudio Estevez, and Samuel Montejo-Sánchez. 2023. A Novel Approach of a Low-Cost Voltage Fault Injection Method for Resource-Constrained IoT Devices: Design and Analysis. Sensors 23, 16 (2023), 7180.
- [28] Xhani Marvin Saß, Richard Mitev, Ahmad-Reza Sadeghi, and Voltage Fault Injection VFI. 2023. Oops..! I Glitched It Again! How to Multi-Glitch the Glitching-Protections on ARM TrustZone-M. arXiv preprint arXiv:2302.06932 (2023).
- [29] Carlton Shepherd, Konstantinos Markantonakis, Nico van Heijningen, Driss Aboulkassimi, Clément Gaine, Thibaut Heckmann, and David Naccache. 2021. Physical fault injection and side-channel attacks on mobile devices: A comprehensive analysis. *Computers & Security* 111 (2021), 102471. https://doi.org/10. 1016/j.cose.2021.102471
- [30] Adrian Tang, Simha Sethumadhavan, and Salvatore Stolfo. 2017. CLKSCREW: Exposing the Perils of Security-Oblivious Energy Management. In 26th USENIX Security Symposium (USENIX Security 17). USENIX Association, Vancouver, BC, 1057–1074. https://www.usenix.org/conference/usenixsecurity17/technicalsessions/presentation/tang
- [31] The U-Boot development community. 2021. Android Verified Boot 2.0 Das U-Boot unknown version documentation. https://docs.u-boot.org/en/v2021. 04/android/avb2.html?highlight=rpmb#avb-using-op-tee-optional [Online; accessed 4. Feb. 2024].
- [32] Jasper van Woudenberg and Colin O'Flynn. 2022. The Hardware Hacking Handbook (1 ed.). No starch press, San Francisco, CA, Chapter Chapter 14: Think of the Children: Countermeasures, Certifications and Goodbytes.
- [33] Vincent Werner, Laurent Maingault, and Marie-Laure Potet. 2023. An end-to-end approach to identify and exploit multi-fault injection vulnerabilities on microcontrollers. *Journal of Cryptographic Engineering* 13, 2 (2023), 149–165.
- [34] Western Digital. 2020. Replay Protected Memory Block (RPMB) Protocol Vulnerabilities. White Paper.
- [35] Z3X-Team. 2024. EASY-JTAG PLUS ACTIVATION. https://z3x-team.com/ products/easy-jtag-plus-activation/ [Online; accessed 5. Feb. 2024].
- [36] Zuluscsi. 2024. ZuluSCSI-firmware. https://github.com/ZuluSCSI/ZuluSCSIfirmware/blob/main/lib/ZuluSCSI_platform_RP2040/sdio_RP2040.pio [Online; accessed 4. Feb. 2024].