

Mind the Gap –
Uncovering the Android patch gap through binary-only patch analysis
HITB conference, April 13, 2018

Jakob Lell <jakob@srlabs.de>
Karsten Nohl <nohl@srlabs.de>



Security
Research
Labs

Allow us to take you on two intertwined journeys

This talk in a nutshell


Research journey

- Wanted to understand how fully-maintained Android phones can be exploited
- Found surprisingly large patch gaps for many Android vendors
- Also found Android exploitation to be unexpectedly difficult



Engineering journey

- Wanted to check thousands of firmwares for the presence of hundreds of patches
- Developed and scaled a rather unique analysis method
- Created an app for your own analysis



Android patching is a known-hard problem

Patching is hard to start with


Patching challenges

- Computer OS vendors regularly issue patches
- Users “only” have to confirm the installation of these patches
- Still, enterprises consider regular patching among the most effortful security tasks

The nature of Android makes patching so much more difficult

- “The mobile ecosystem’s diversity [...] contributes to security update complexity and inconsistency.” – FTC report, March 2018 ^[1]
- Patches are handed down a long chain of typically four parties before reaching the user
- Only some devices get patched (2016: 17% ^[2]). We focus our research on these “fully patched” phones

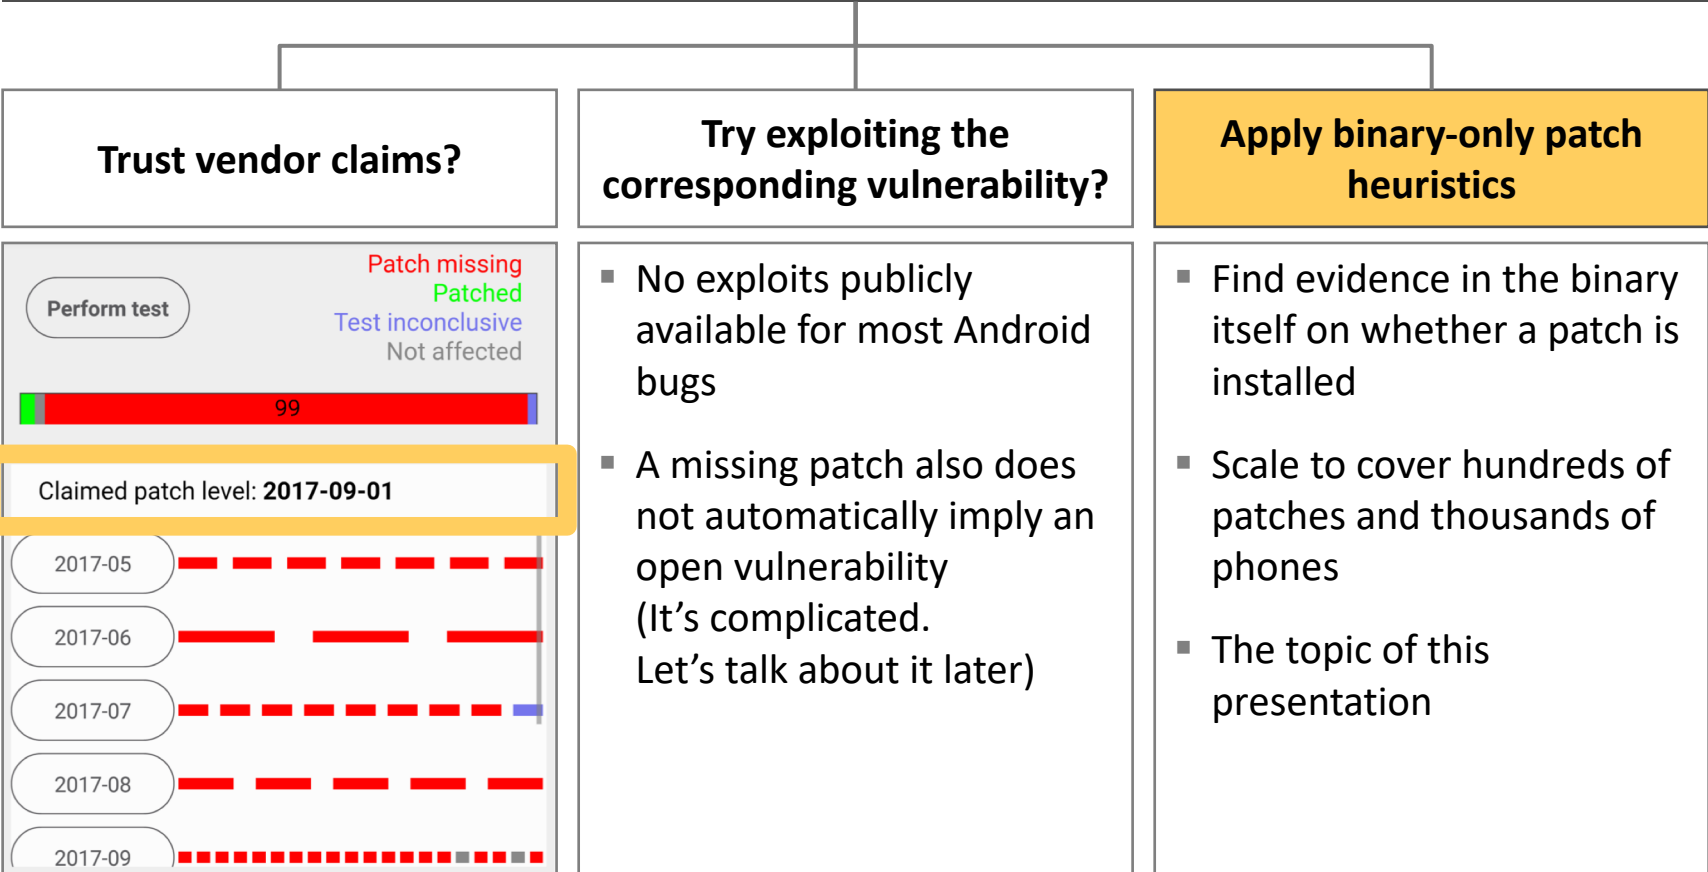
Patch ecosystems



Our research question – How many patching mistakes are made in this complex Android ecosystem? That is: how many patches go missing?


Vendor patch claims can be unreliable; independent verification is needed

How do we determine whether an Android binary has a patch installed, without access to the corresponding source code?



Important distinction: A missing **patch** is *not* automatically an open security **vulnerability**. We'll discuss this a bit later.

Patching is necessary in the Android OS and the underlying Linux kernel

	Android OS patching (“userland”)	Linux kernel patching
Responsibility	<ul style="list-style-type: none">▪ Android Open Source Project (AOSP) is maintained by Google▪ In addition, chipset and phone vendors extend the OS to their needs	<ul style="list-style-type: none">▪ Same kernel that is used for much of the Internet▪ Maintained by a large ecosystem▪ Chipset and phone vendors contribute hardware drivers, which are sometimes kept closed-source
Security relevance	<ul style="list-style-type: none">▪ Most exposed attack surface: The OS is the primary layer of defense for remote exploitation	<ul style="list-style-type: none">▪ Attackable mostly from within device▪ Relevant primarily for privilege escalation (“rooting”)
Patch situation	<ul style="list-style-type: none">▪ Monthly security bulletins published by Google▪ Clear versioning around Android, including a patch level date, which Google certifies for some phones	<ul style="list-style-type: none">▪ Large number of vulnerability reports, only some of which are relevant for Android▪ Tendency to use old kernels even with latest Android version; e.g., Kernel 3.18 from 2014, end-of-life: 2017
We focus our attention on userland patches 		

Agenda

-
- Research motivation

 **Spot the Android patch gap**

- Try to exploit Android phones
-

We want to check hundreds of patches on thousands of Android devices

Android userland patch analysis

Android's 2017 security bulletins list
~280
bugs (~CVEs) with Critical or High severity

Source code is available for
~240
of these bugs

Of these userland bugs,
~180
originate from C/C++ code (plus a few Java)

So far, we implemented heuristics for
164
of the corresponding patches



The heuristics would optimally work on hundreds of thousands of Android firmwares:
– 60,000 Android variants^[3]
– Regular updates for many of these variants

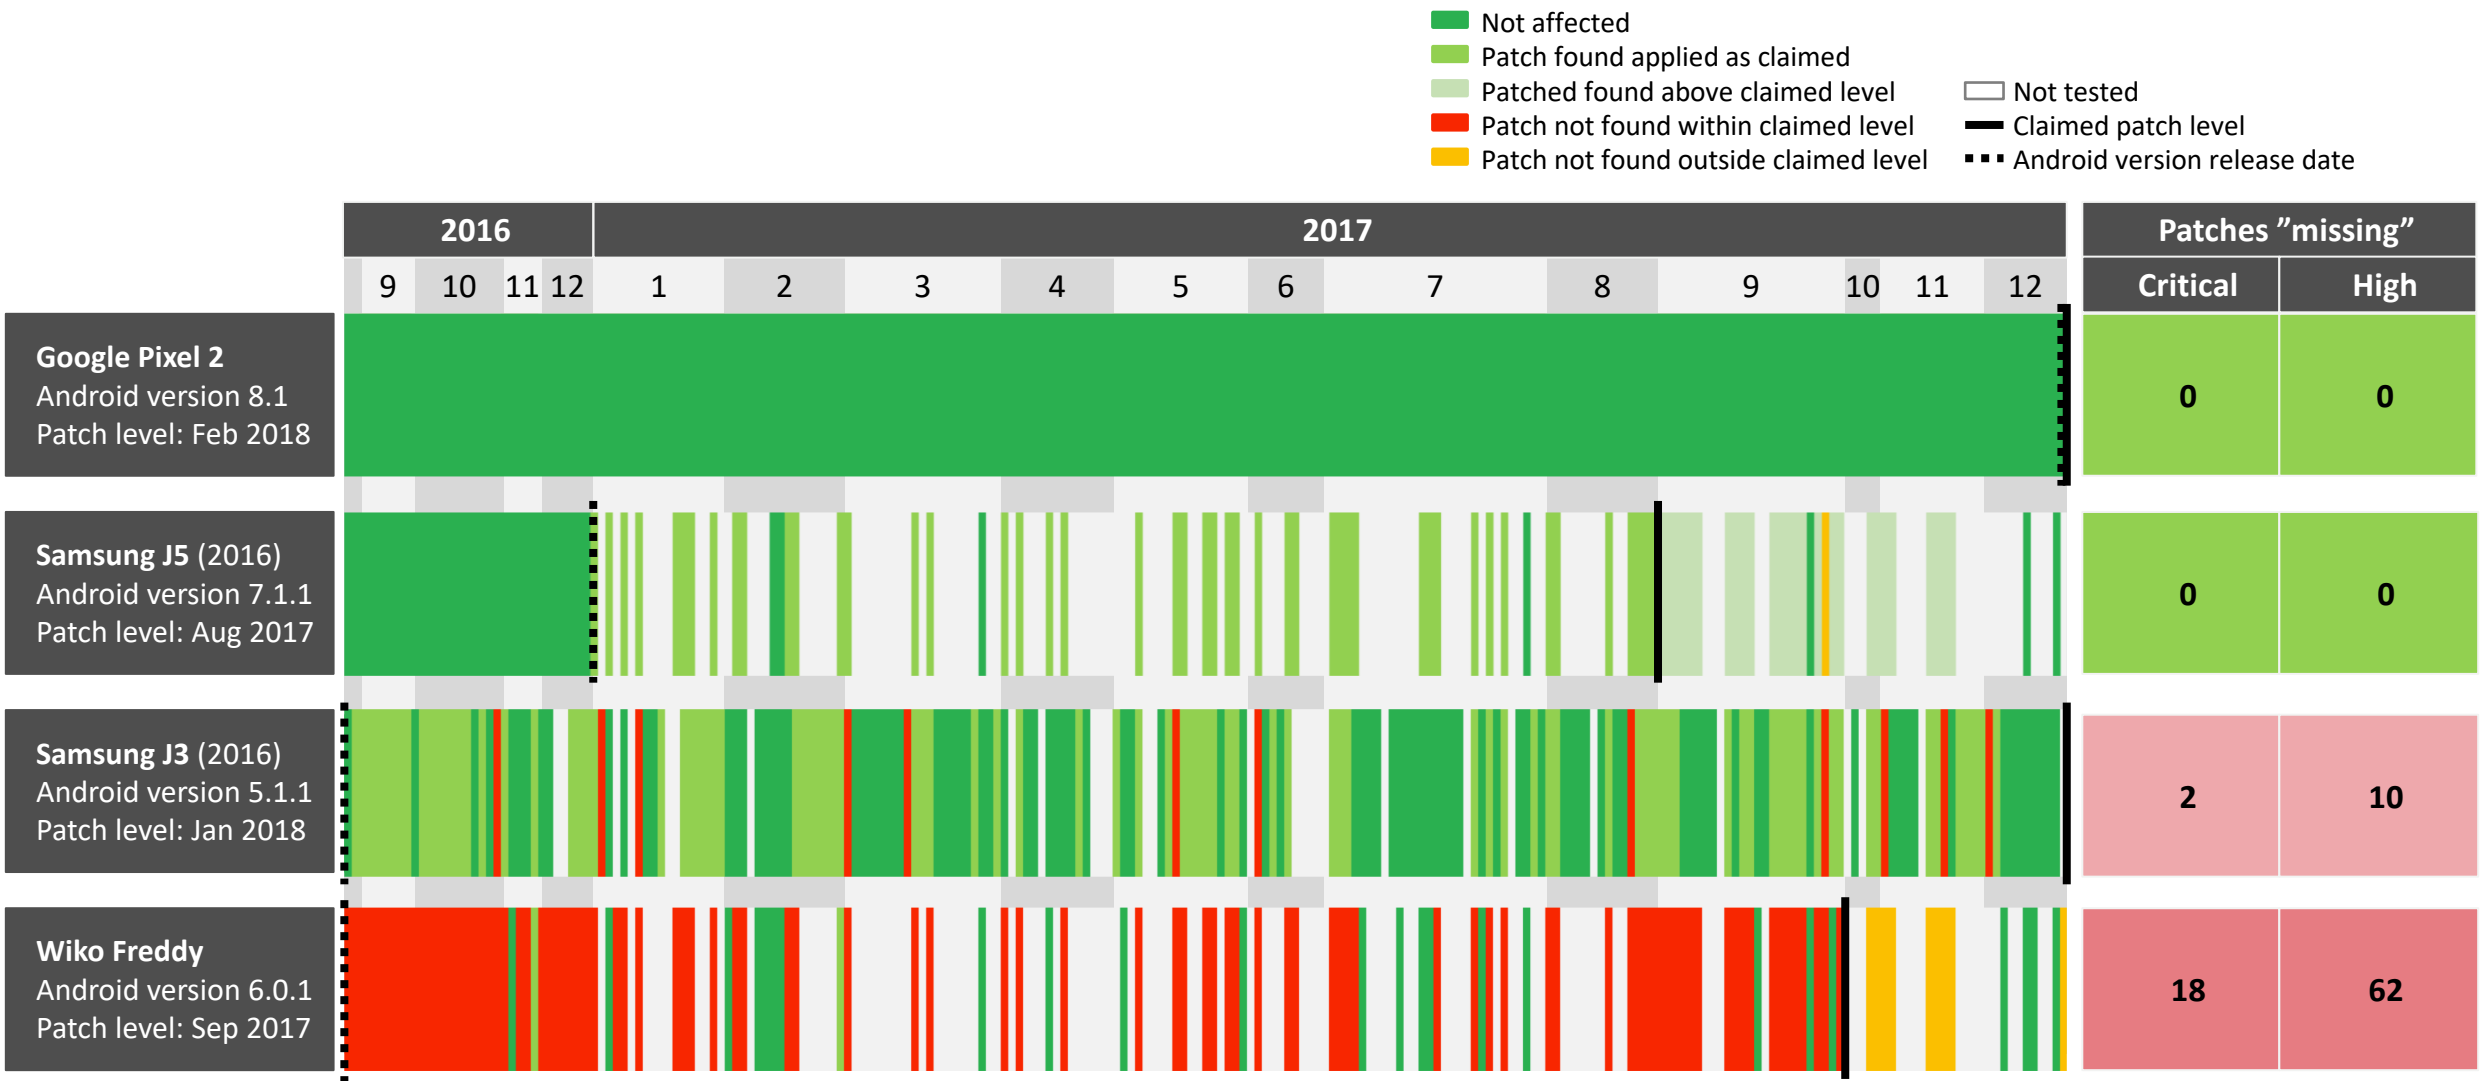
Out-of-scope (for now)

~700 kernel and medium/low severity userland patches

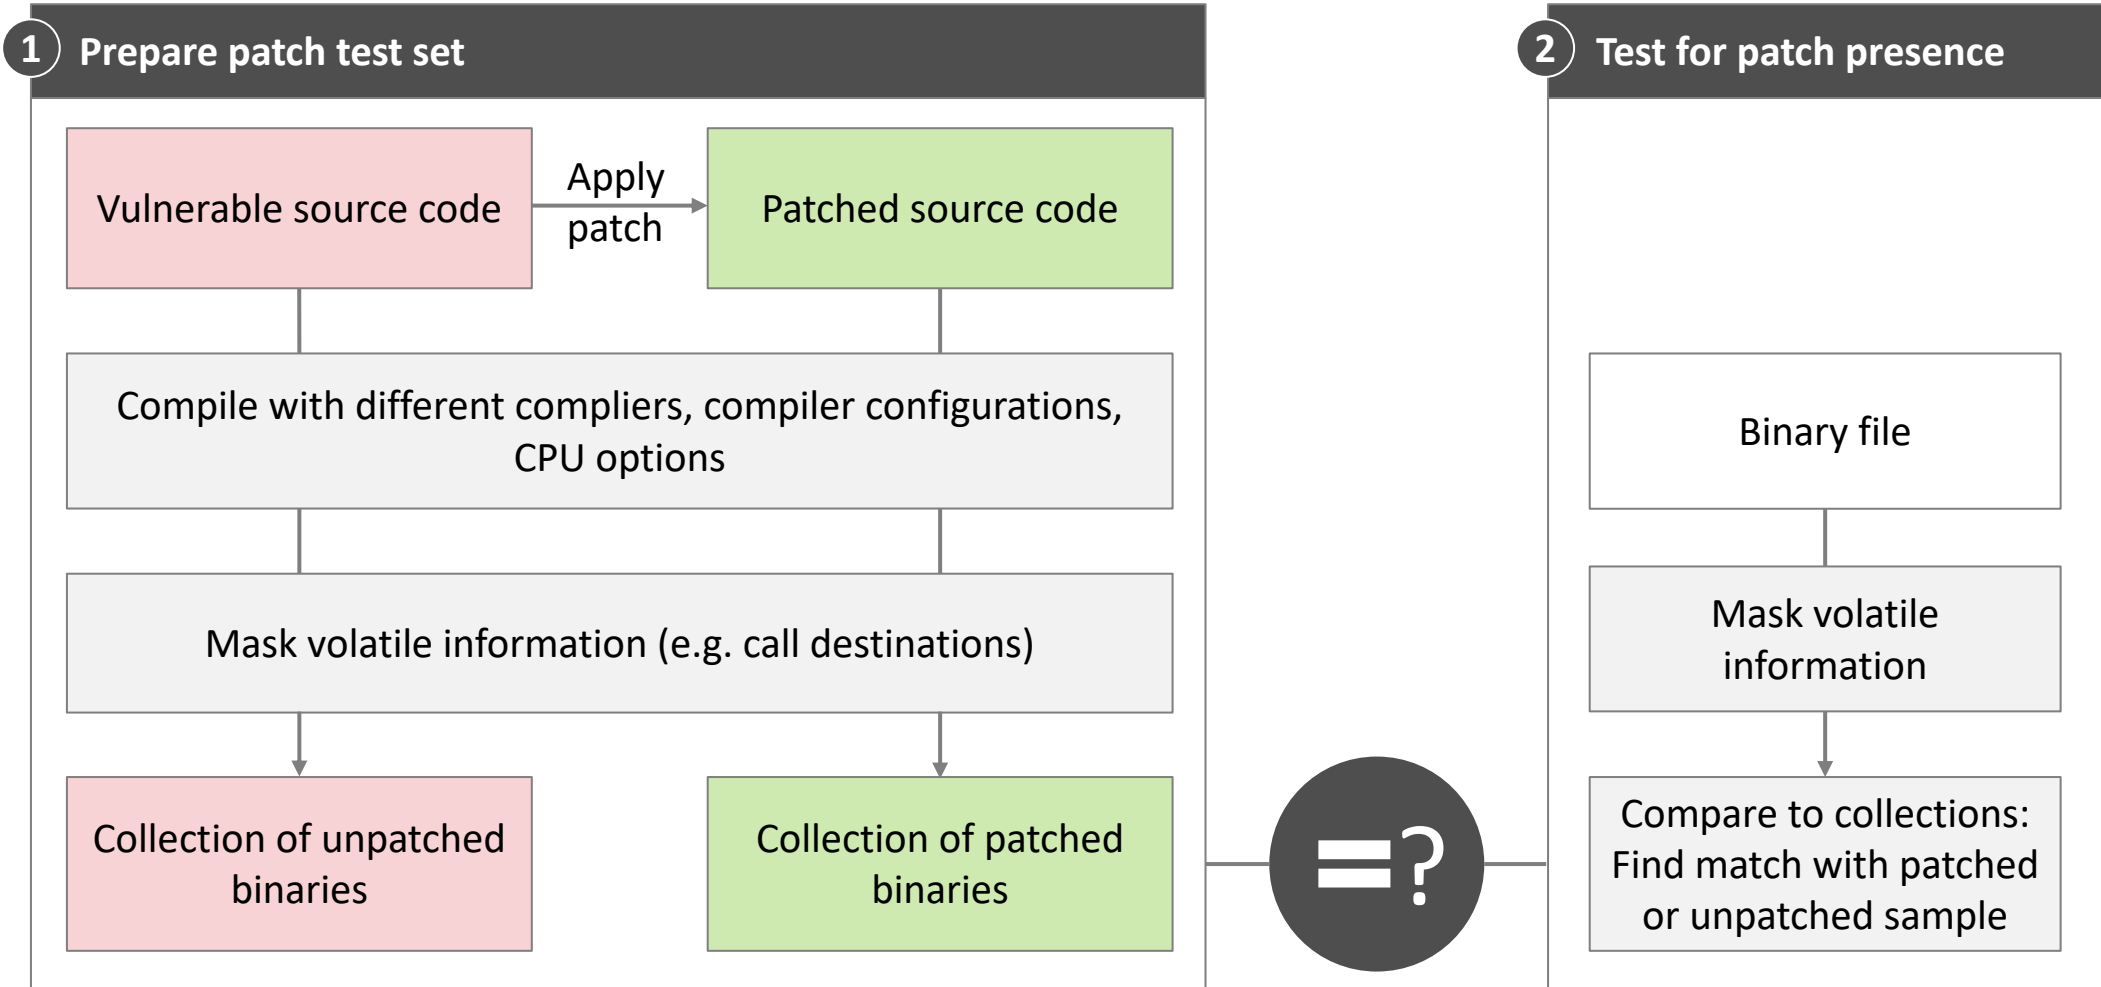
The remaining bugs are in closed-source vendor-specific components

We do not yet support most Java patches

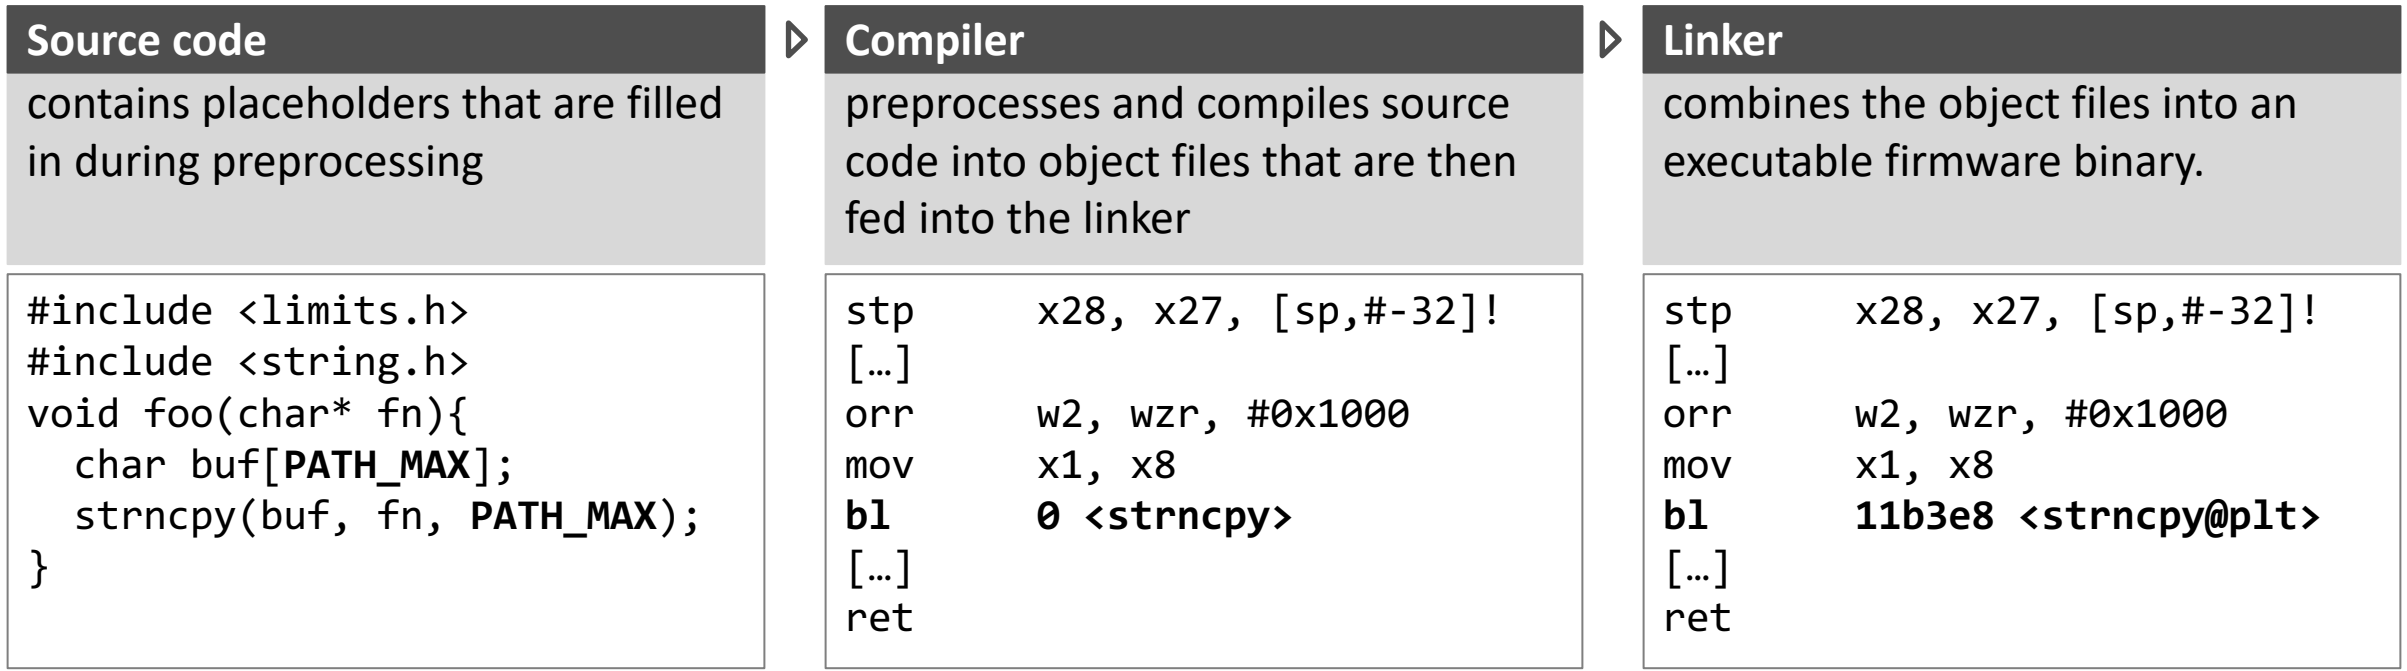
The patch gap: Android patching completeness varies widely for different phones



Binary-only analysis: Conceptually simple



A bit more background: Android firmwares go from source code to binaries in two steps



The basic idea: Signatures can be generated from reference source code

1

Compile reference source code (before and after patch)

Parse disassembly listing for relocation entries

```
Disassembly of object file, after compiler but before linker
0000000000000000 <impeg2d_api_reset>:
  0: a9bd7bfd stp    x29, x30, [sp, #-48]!
  4: 910003fd mov    x29, sp
[...]
 20: f9413e60 ldr    x0, [x19, #632]
 24: 52800042 mov    w2, #0x2 // #2
 28: b9402021 ldr    w1, [x1, #32]
2c: 94000000 bl     0 <impeg2_buf_mgr_release> 2c: R_AARCH64_CALL26 impeg2_buf_mgr_release
[...]
```

Prepare patch test set

Sanitize instructions
Toss out irrelevant destination addresses of the instruction



Create hash of remaining binary code

Generate signature containing function length, position/type of relocation entries, and hash of the code

At scale, three compounding challenges need to be solved

Too much source code

- There is too much source code to collect
- Once collected, there is too much source code to compile



Too many compilation possibilities

- Hard to guess which compiler options to use
- Need to compile same source many times



Hard to find code “needles” in binary “haystacks”

- Without symbol table, whole binary needs to be scanned
- Thousands of signatures of arbitrary length

Signature generation would require huge amounts of source code

One Android source code tree is roughly 50 GiB in size

Source code trees are managed in a manifest, which lists git repositories with revision and path in a source code tree

```
...  
<project name="platform/external/zxing" revision="d2256df36df8778a3743e0a71eab0cc5106b98c9"/>  
<project name="platform/frameworks/av" revision="330d132dfab2427e940cfaf2184a2e549579445d"/>  
<project name="platform/frameworks/base" revision="85838feaea8c8c8d38c4262e74d911e59a275d02"/>  
...  
+~500 MORE REPOSITORIES
```

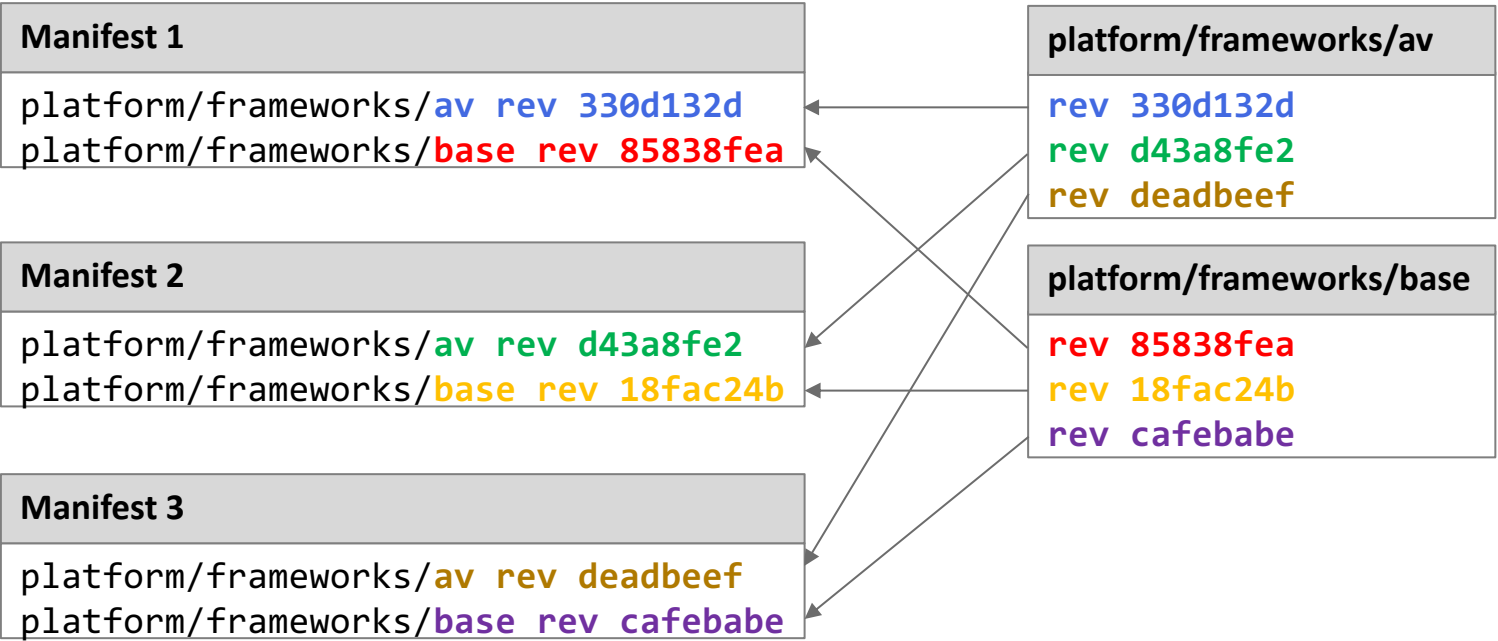
Signature generation requires many source code trees

- **Hundreds of different Android revisions**
(e.g. android-7.1.2_r33)
- **Device-specific source code trees**
(From Qualcomm Codeaurora CAF)

Currently ~1100 source code trees are used in total
(many more exist!)
1100 x 50 GiB = 55 TiB
Would require huge amount of storage, CPU time, and network traffic to check out everything

We leverage a FUSE (filesystem in userspace) to retrieve files only on demand

Insight: The same git repositories are used for many manifests.



How this can be leveraged

- Filesystem in userspace (FUSE)**
- Store each git repository only once (with `git clone --no-checkout`)
 - Extract files from git repository on demand when the file is read
 - Use database for caching directory contents

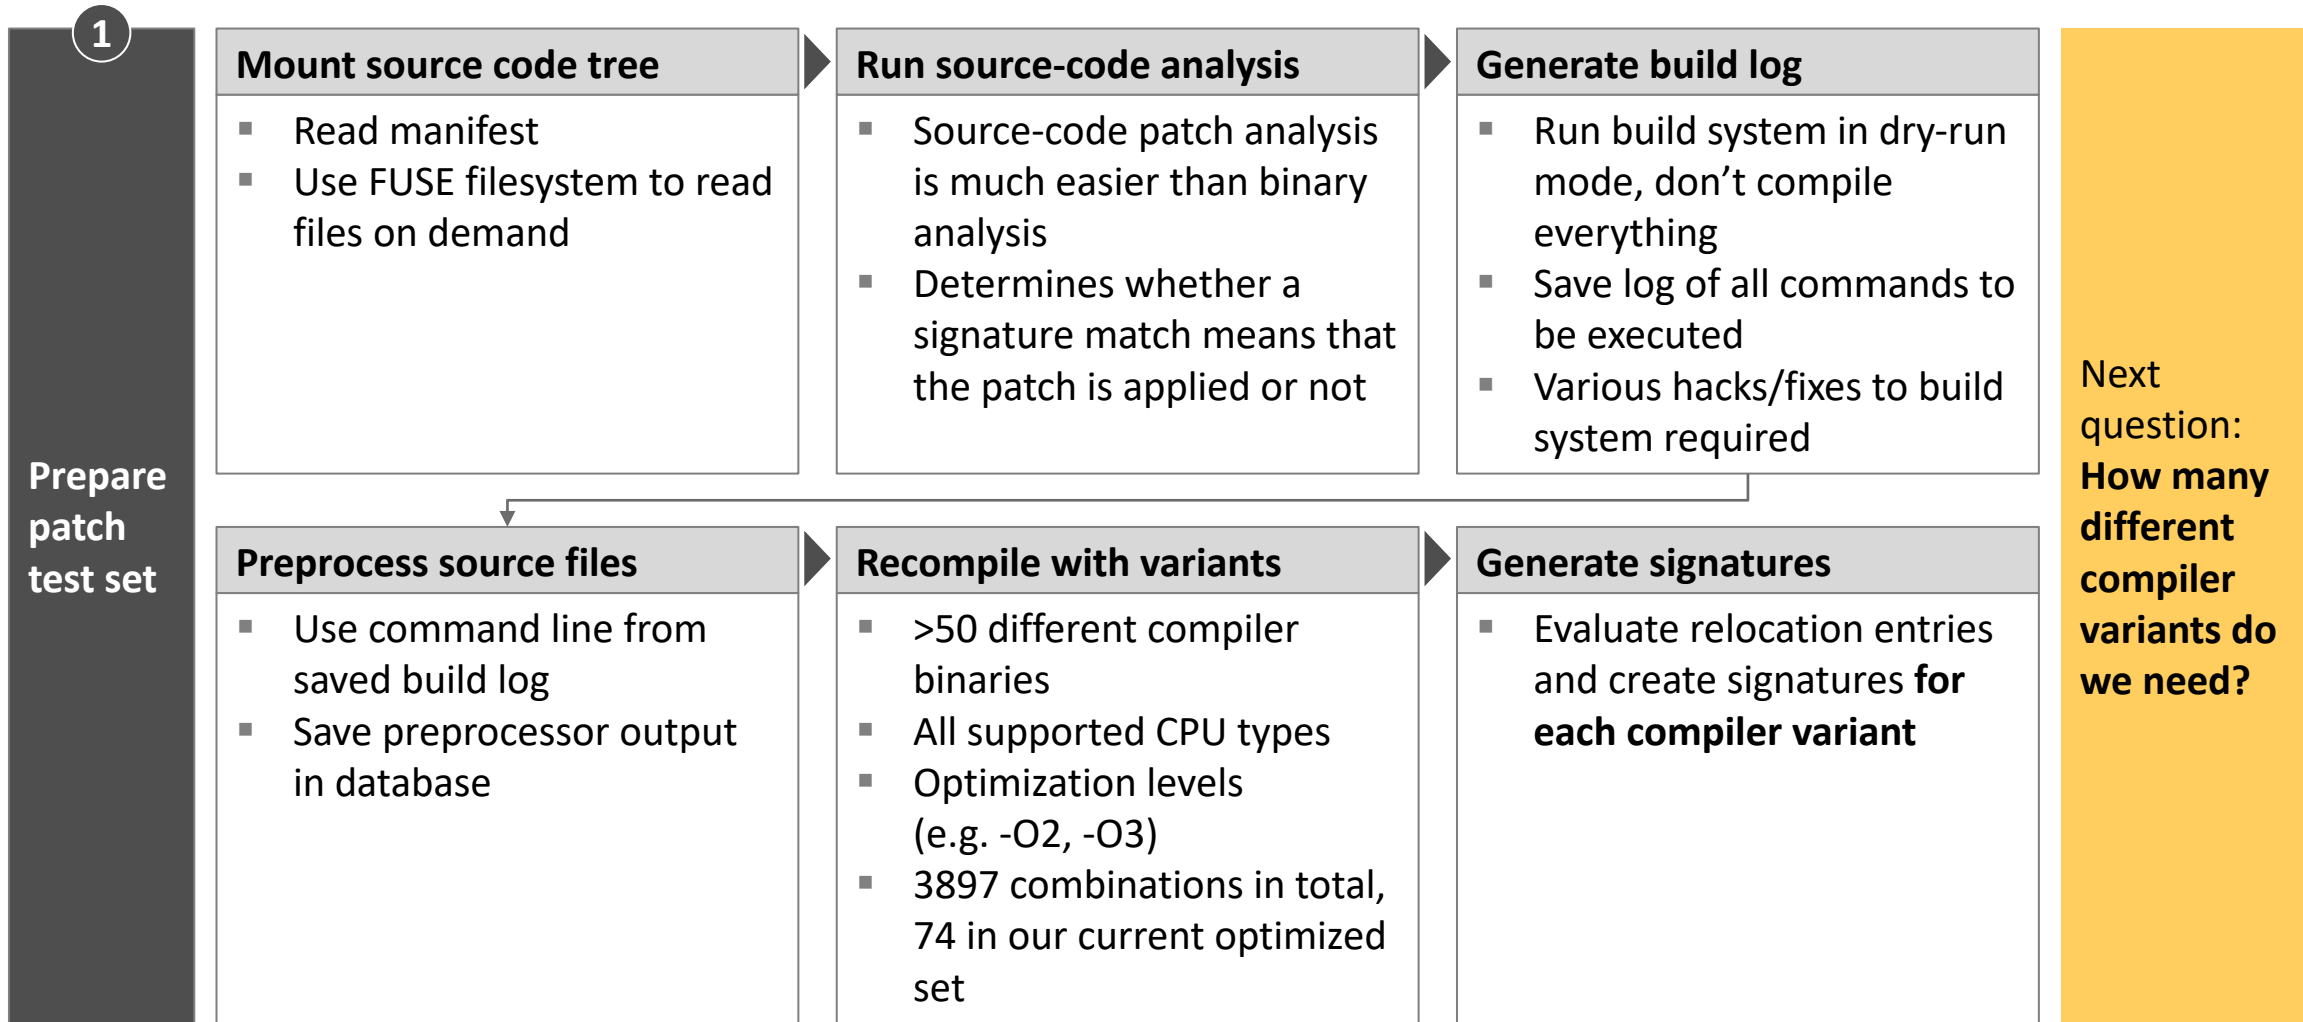
**Reduces storage requirement by >99%:
55 TiB => 300 GiB**

Saves network bandwidth and time required for checkout

Prevents IP blocking by repository servers

Using our custom FUSE, we can finally generate a large collection of signatures

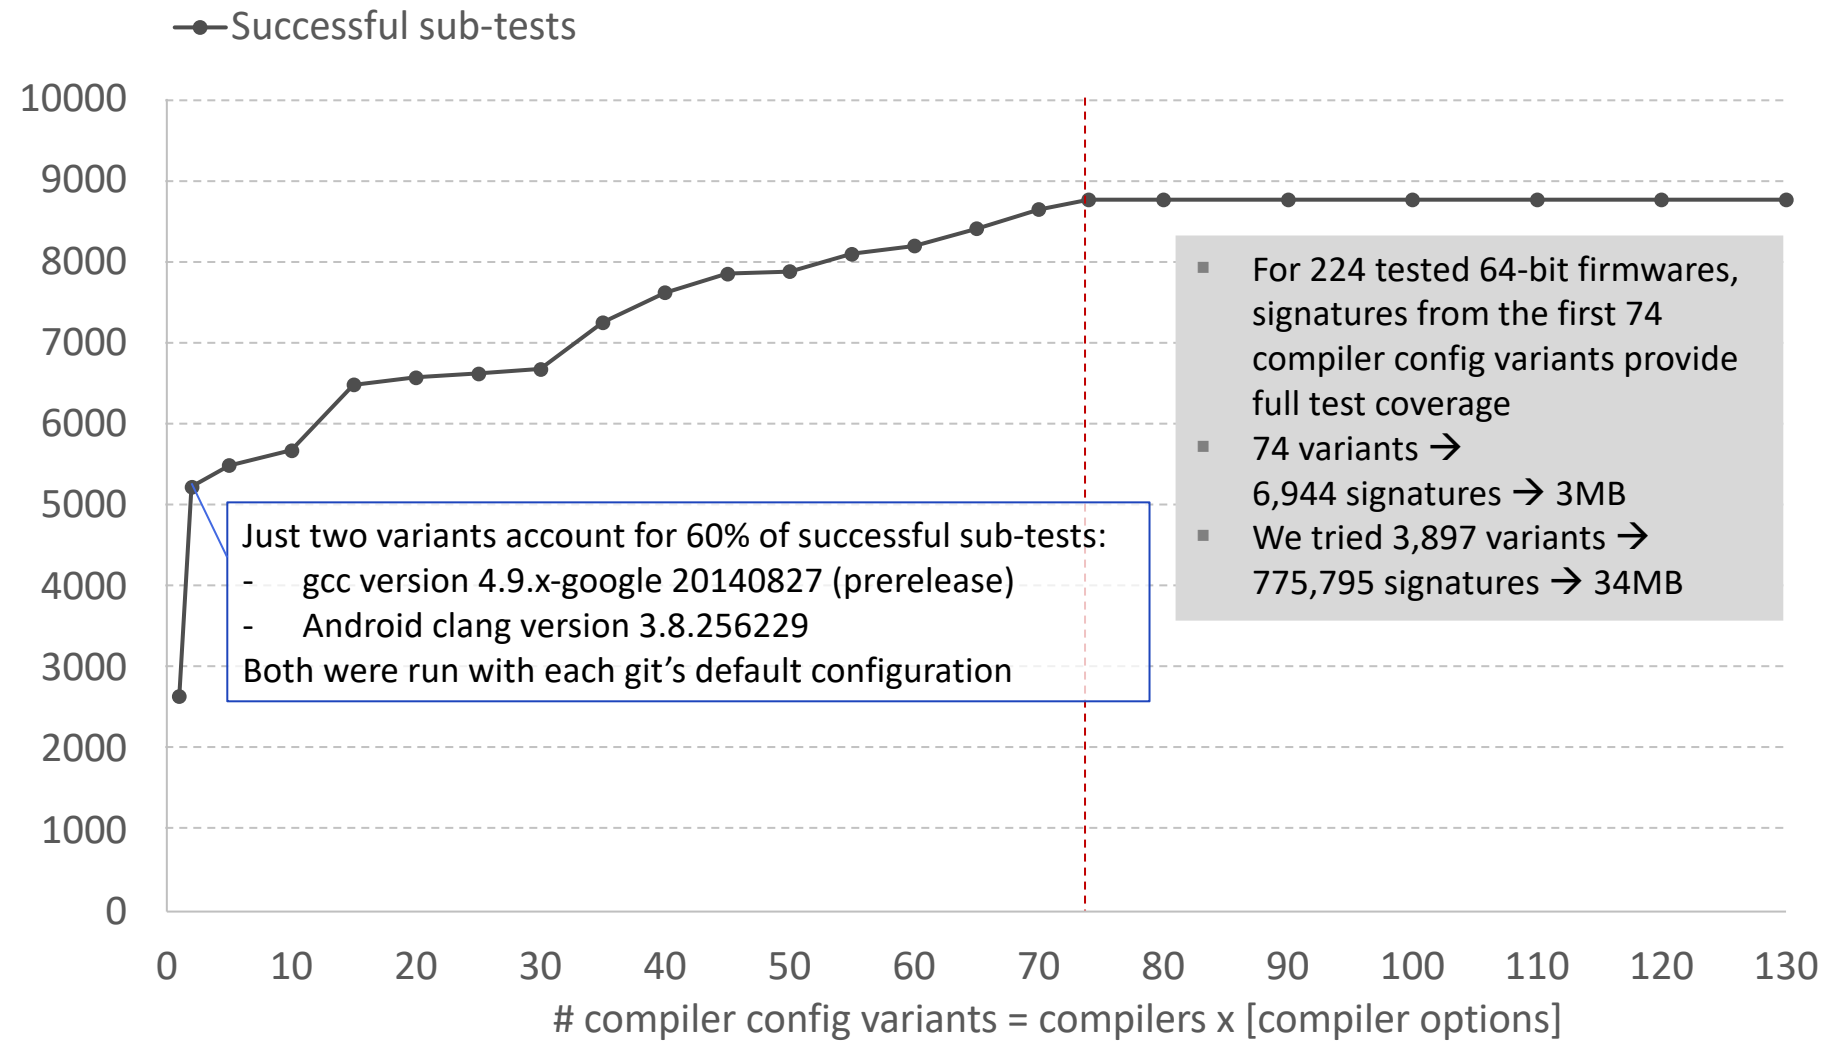
- ▶ Amount of source code
- ▶ Compilation possibilities
- ▶ Needles in haystacks



Brute-forcing 1000s of compiler variants finds 74 that produce valid signatures for all firmwares tested to date

Tests are regularly optimized

- Our collection includes 3897 compiler configuration variants, only 74 of which are required for firmwares tested to date.
- To ensure a high rate of conclusive tests, test results are regularly checked for success.
- The test suite is amended with additional variants from the collection as needed.
- The collection itself is amended with additional compiler configuration variants as they become relevant.



- For 224 tested 64-bit firmwares, signatures from the first 74 compiler config variants provide full test coverage
- 74 variants → 6,944 signatures → 3MB
- We tried 3,897 variants → 775,795 signatures → 34MB

Just two variants account for 60% of successful sub-tests:
- gcc version 4.9.x-google 20140827 (prerelease)
- Android clang version 3.8.256229
Both were run with each git's default configuration

Finding needles in a haystack: What do we do if there is no symbol table?

2

Function found in symbol table

Simply compare function with pre-computed samples

Function not in symbol table

Challenge

Insight

Solution

Checking signature at each position is computationally expensive

Similar problem already solved by rsync

Take advantage of rsync rolling checksum algorithm

Relocation entries are not known while calculating checksum

Relocation entries are only used for certain instructions

Guess potential relocation entries based on instruction type and sanitize args before checksumming

32bit code uses Thumb encoding, for which instruction start is not always clear

Same binary code is often also available in 64bit version based on same source code

Only test 64bit code

Test for patch presence

Using improved rolling signatures, we can efficiently search the binary 'haystack' for our code 'needles'

	Process step	Hex dump of instruction	Assembly code / instructions
Sanitize arguments before checksumming	Potential relocation entries are detected based on instruction.	...	
	Zero-out volatile bits	97fee7a2 94000000 f10002ff 1a9f17e8 b40000b6 3707fdc8 f10006d6 54ffff42 35fffd48 36000255 394082e8 35000208 52adad21 320003e8 728daca1	b1 c7c40 <strncpy@plt> b1 0 cmp x23, #0x0 cset w8, eq cbz x22, 10ddbc tbnz w8, #0, 10dd6 subs x22, x22, #0x1 b.cs 10dd9c cbnz w8, 10dd64 tbz w21, #0, 10de08 ldrb w8, [x23,#32] cbnz w8, 10de08 mov w1, #0x6d690000 orr w8, wzr, #0x1 movk w1, #0x6d65
Match signatures of arbitrary lengths using sliding windows	Size-8 window matches on start of signature		
	Overlapping window matches on end of signature		

To avoid false positives (due to guessed relocation entries), signature is matched from the first window to the end of the overlapping window

Putting it all together: With all three scaling challenges overcome, we can start testing

1

Prepare patch test set

Mount source code tree

- Read manifest
- Fuse filesystem to read files on demand

Run source-code analysis

- Source-code patch analysis is much easier than binary analysis
- Determines whether a signature match means that the patch is applied or not

Generate build log

- Run build system in dry-run mode, don't compile everything
- Save log of all commands to be executed
- Various hacks/fixes to build system required

Preprocess source files

- Use command line from saved build log
- Save preprocessor output in database

Recompile with variants

- >50 different compiler binaries
- All supported CPU types
- Optimization levels (e.g. -O2, -O3)
- 3897 combinations in total, 74 in our current optimized set

Generate signatures

- Evaluate relocation entries and create signatures **for each compiler variant**

2

Test for patch presence

- Find and extract function (using symbol table or rolling signature)
- Mask relocation entries from signature
- Calculate and compare hash of remaining code

Patch gap: Android vendors differ widely in their patch completeness

Vendors differ in how many patches are missing from their phones	Missed patches	Vendor	Samples*	Notes
	0 to 1	Google	Lots	<ul style="list-style-type: none"> The tables shows the average number of missing Critical and High severity patches before the claimed patch date * Samples – Few: 5-9; Many: 10-49; Lots: 50+ Some phones are included multiple times with different firmwares releases Not all patch tests are always conclusive, so the real number of missing patches could be higher Not all patches are included in our tests, so the real number could be higher still Only phones are considered that were patched October-2017 or later A missing patch does not automatically indicate that a related vulnerability can be exploited
		Sony	Few	
		Samsung	Lots	
		Wiko	Few	
	1 to 3	Xiaomi	Many	
		OnePlus	Many	
		Nokia	Few	
	3 to 4	HTC	Few	
		Huawei	Many	
LG		Many		
Motorola		Many		
More than 4	TCL	Many		
	ZTE	Few		

Some of the patch gap is likely due to chipset vendors forgetting to include them	Missed patches	Chipset	Samples*	Notes
	< 0.5	Samsung	Lots	<ul style="list-style-type: none"> Again, we show the average of missing High and Critical patches for phones that use these chipsets Samsung phones can run on a Samsung or Qualcomm chipset
	1.1	Qualcomm	Lots	
	1.9	HiSilicon	Many	
	9.7	Mediatek	Many	

Agenda

-
- Research motivation
 - Spot the Android patch gap

 **Try to exploit Android phones**

Can we now hack Android phones due to missing patches?

At first glance, Android phones look hackable

- We find that most phones miss patches within their patch level
- While the number of open CVEs can be smaller than the number of missing patches, we expect some vulnerabilities to be open
- Many CVEs talk of “code execution”, suggesting a hacking risk based on what we experience on Windows computers

VS.

Mobile operating systems are inherently difficult to exploit

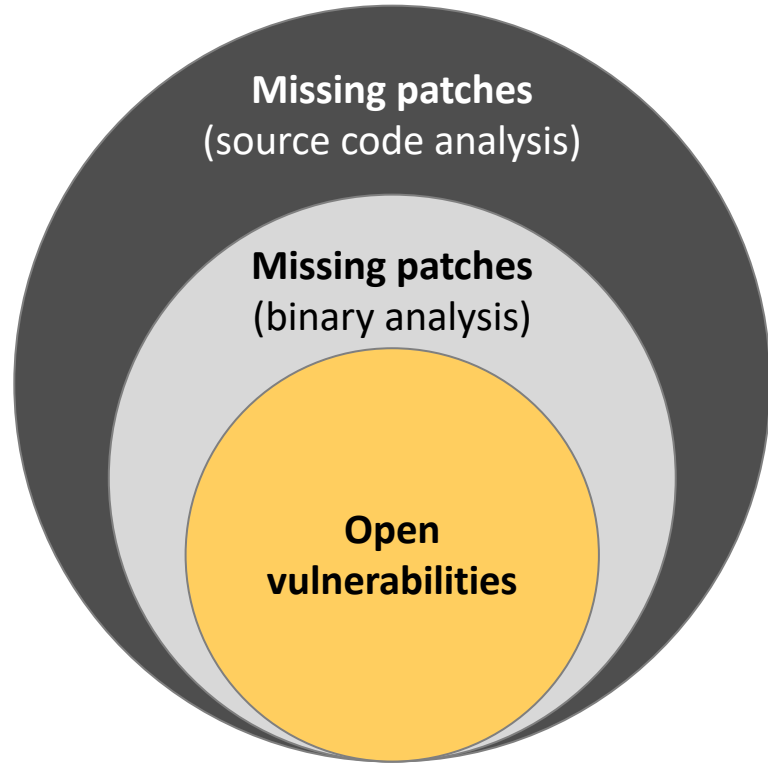
- Modern exploit mitigation techniques increase hacking effort
- Mobile OSs explicitly distrust applications through sandboxing, creating a second layer of defense
- Bug bounties and Pwn2Own offer relatively high bounties for full Android exploitation

Do criminals hack Android? Very rarely.

Criminals generally use three different methods to compromise Android devices			
	Social engineering	Local privilege escalation	Remote compromise
Approach	<p>Trick user into insecure actions:</p> <ul style="list-style-type: none"> Install malicious app Then grant permissions Possibly request 'device administrator' role to hinder uninstallation 	<ul style="list-style-type: none"> Trick user into installing malicious app Then exploit kernel-level vulnerability to gain control over device, often using standard "rooting" tools 	<ul style="list-style-type: none"> Exploit vulnerability in an outside-facing app (messenger, browser) Then use local privilege escalation
Used for	<ul style="list-style-type: none"> Ransomware [File access permission] 2FA hacks [SMS read] Premium SMS fraud [SMS send] 	<ul style="list-style-type: none"> Targeted device compromise, e.g. FinFisher and Crysaor (Same company as infamous Pegasus malware) Advanced malware 	<ul style="list-style-type: none"> (Google bug bounty, Pwn2Own)
Frequency in criminal activity	<ul style="list-style-type: none"> Almost all Android "Infections" ●●● 	<ul style="list-style-type: none"> Regular observed in advanced malware and spying ○○● 	<ul style="list-style-type: none"> Very few examples of recent criminal use ○○○
Made harder through patching	<p>✗</p>	<p>✓ (userland or kernel)</p>	<p>✓ (userland and kernel)</p>

An exploitable vulnerability implies a missing patch, but not the other way around

Missing patches in source code	
—	Code parts that are ignored during compilation
=	Missed patches in binary
—	Vendor created alternative patch
—	Vulnerability requires a specific configuration
—	Bug is simply not exploitable
—	Errors in our heuristic (it happens!)
=	Open vulnerabilities



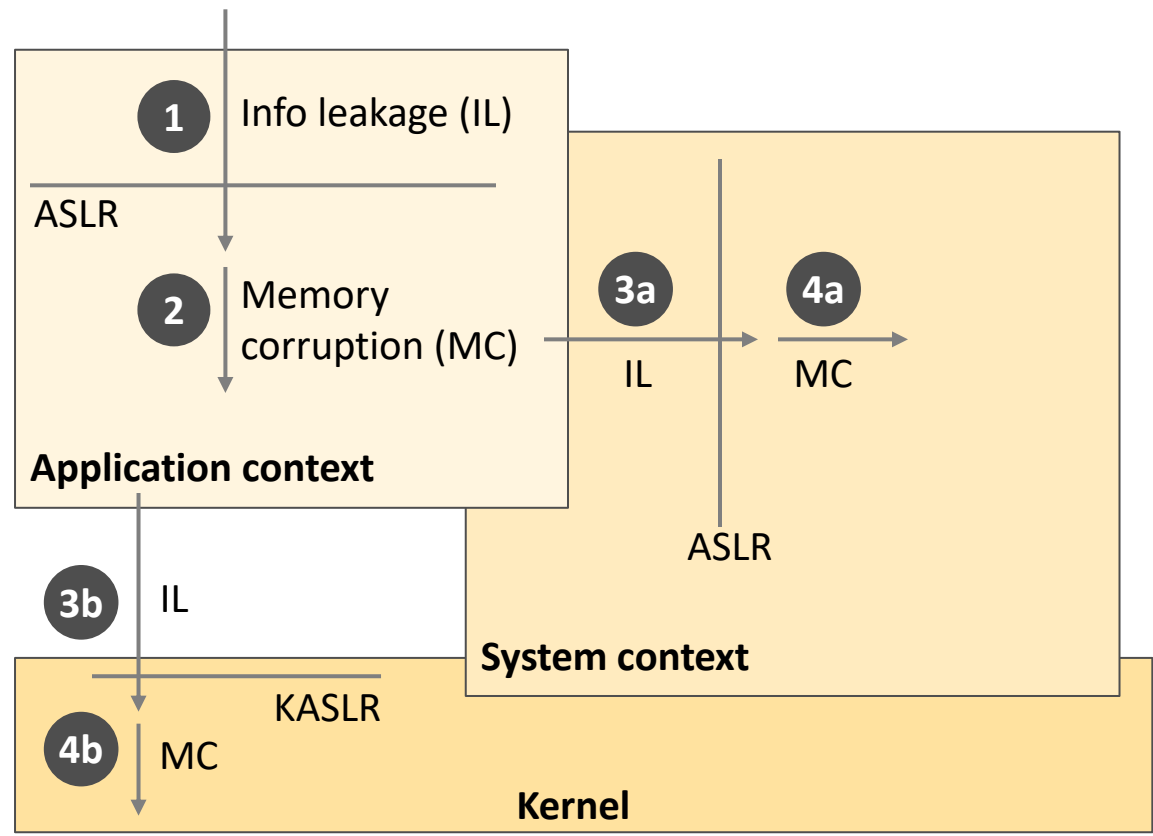
A single Android bug is almost certainly not enough for exploitation

Android remote code execution is a multi-step process

- 1 Information leakage** is used to derive ASLR memory offset (alternatively for 32-bit binaries, this offset can possibly be brute-forces)
- 2 Corrupt memory** in an application. Examples:
 - Malicious video file corrupts memory using Stagefright bug
 - Malicious web site leverages Webkit vulnerability

➤ This gives an attacker control of the application including the apps access permission
- 3 4** Do the same again with two more bugs to gain access to system context or kernel
 - This gives an attacker all possible permissions (system context), or full control over the device (kernel)

Simplified exploit chain examples with 4 bugs

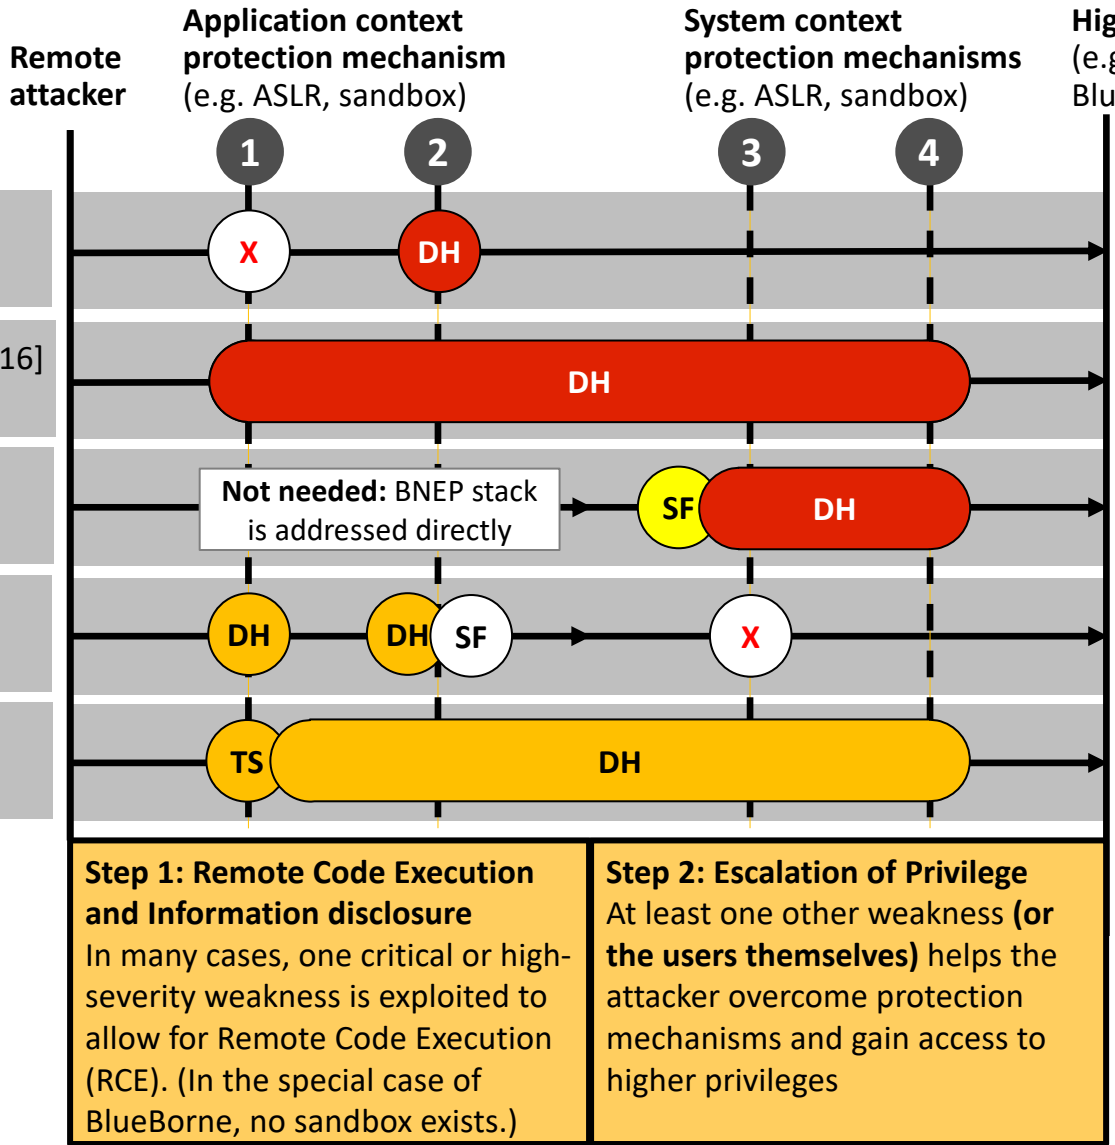


Aside from exploiting MC and IL programming bugs, Android has experienced logic bugs that can enable alternative, often shorter, exploit chains

Remotely hacking a modern Android device usually requires chains of bugs

Famed real-world exploit examples

Stagefright [2015] Android < 5.1.1
Return to libstagefright [2016] Android < 7.0
BlueBorne [2017] Android < 8.0
Pixel - Nexus 6P [2017] Chrome Android prior 54.0.2840.90
Pixel [2018] Chrome Android prior 61.0.3163.79



Step 1: Remote Code Execution and Information disclosure
In many cases, one critical or high-severity weakness is exploited to allow for Remote Code Execution (RCE). (In the special case of BlueBorne, no sandbox exists.)

Step 2: Escalation of Privilege
At least one other weakness (**or the users themselves**) helps the attacker overcome protection mechanisms and gain access to higher privileges

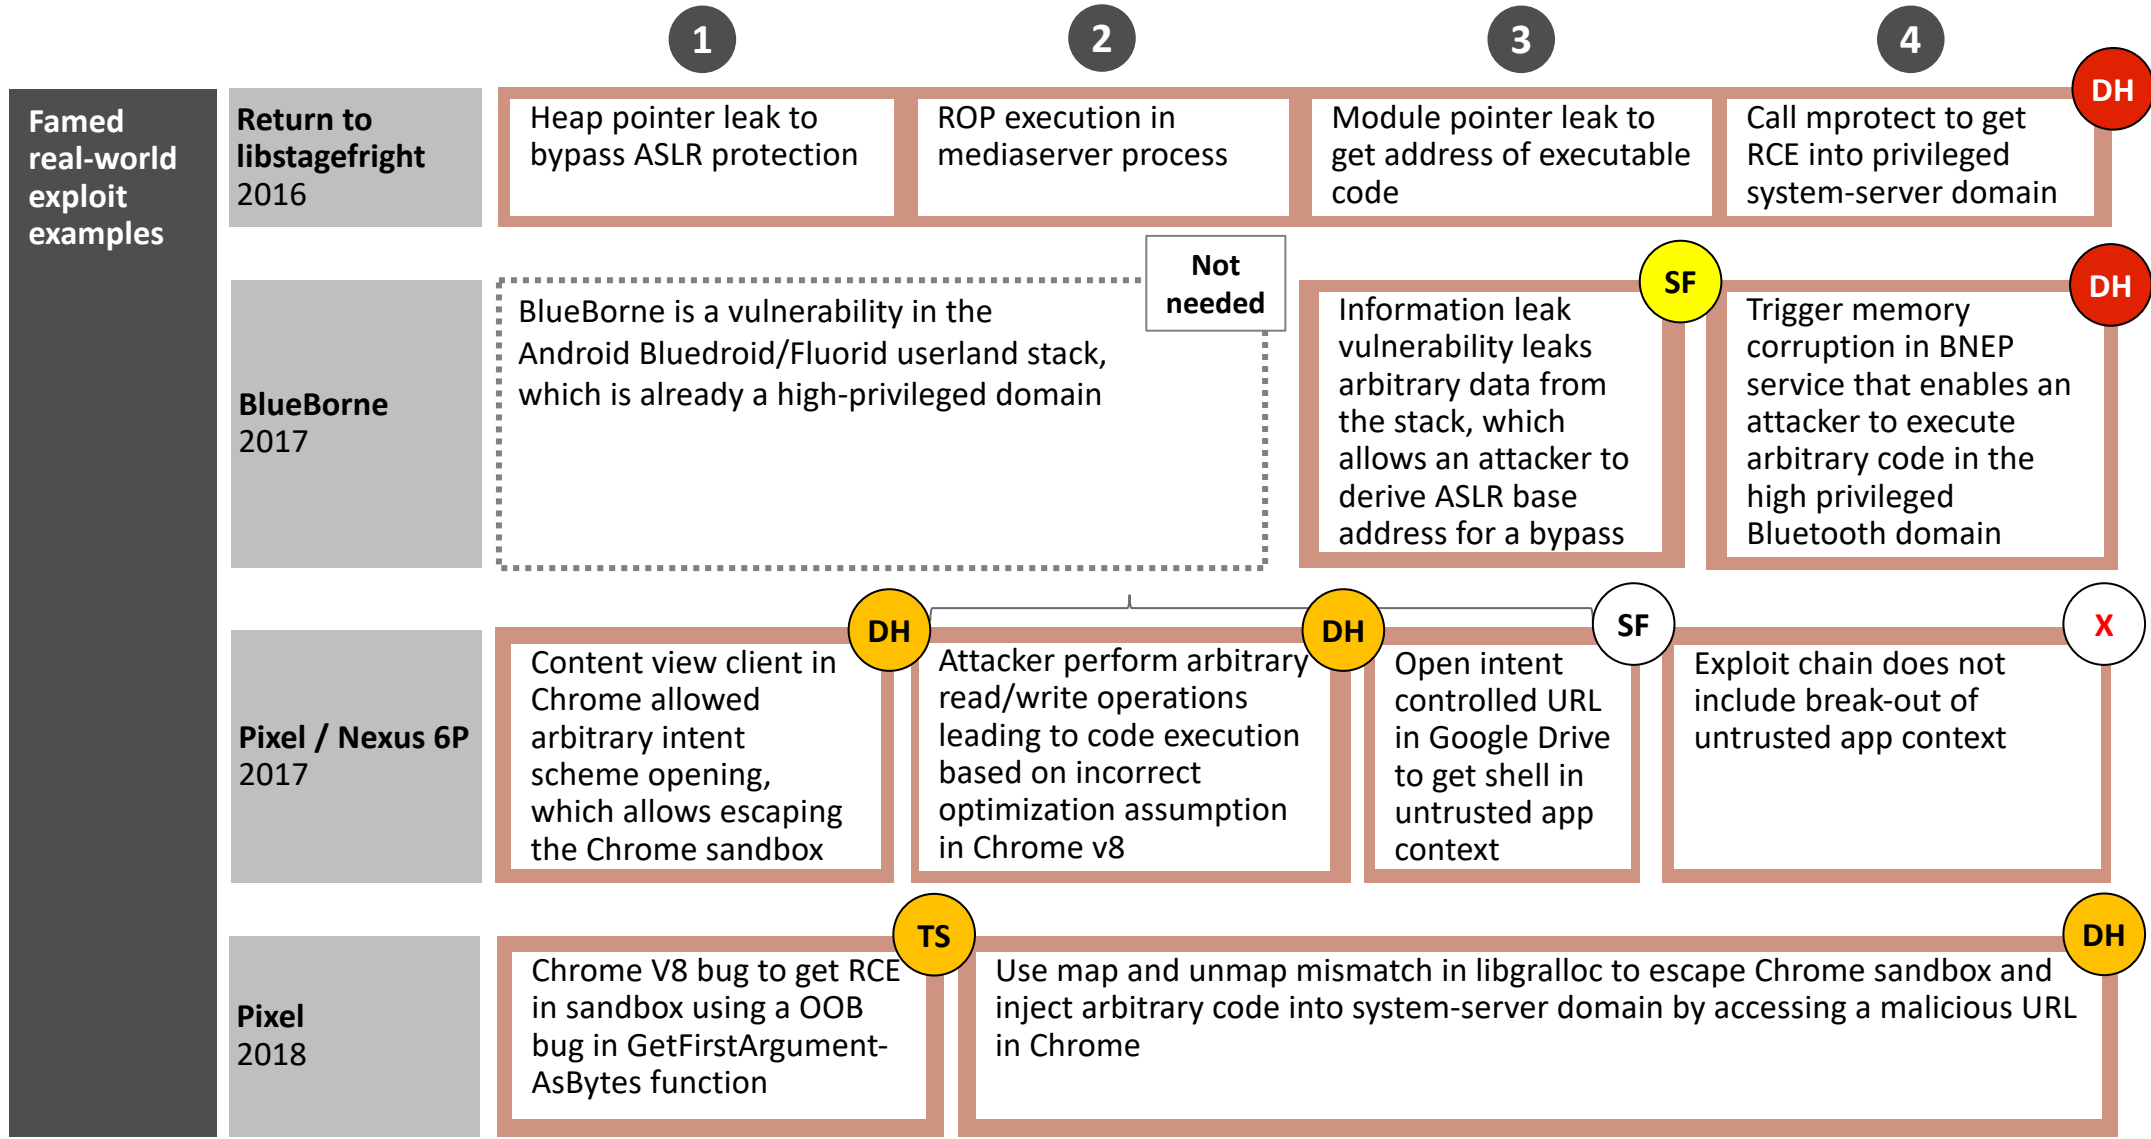
Weakness severities

- Critical
- High
- Moderate

Weakness classes

- DH** Data handling errors (CWE-19)
e.g. buffer errors, input validation mistakes
- SF** Security features gaps (CWE-254)
e.g. permission errors, privileges mishandling, access control errors
- TS** Time and state errors (CWE-361)
e.g. race conditions, incorrect type conversions or casting

In case you want to dive deeper: More details on well-documented Android exploit chains



SnoopSnitch version 2.0 introduces patch analysis for all Android users

Tool name

SnoopSnitch


Purpose

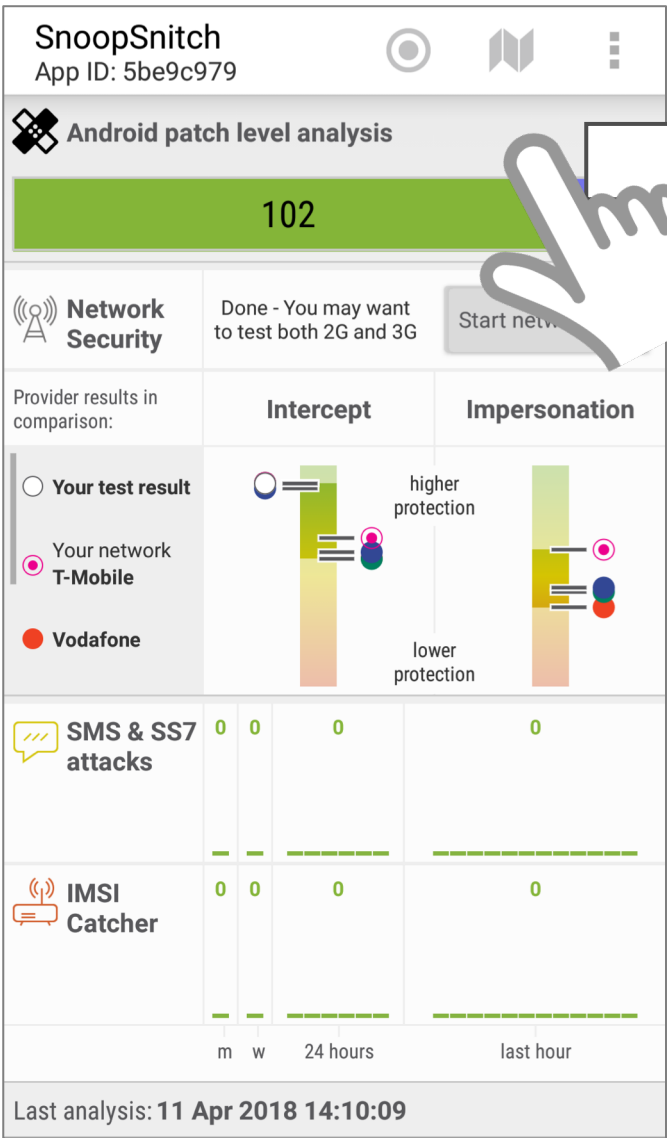
- [new in 2.0] Detect potentially missing Android security patches
- Collect network traces on Android phone and analyze for abuse
- Optionally, upload network traces to GSMmap for further analysis

Requirements

- Android version 5.0
- Patch level analysis: All phones incl. **non-rooted**
- Network attack monitoring: Rooted Qualcomm-based phone

Source


[Search: SnoopSnitch](#)



SnoopSnitch
App ID: 5be9c979

Android patch level analysis

102

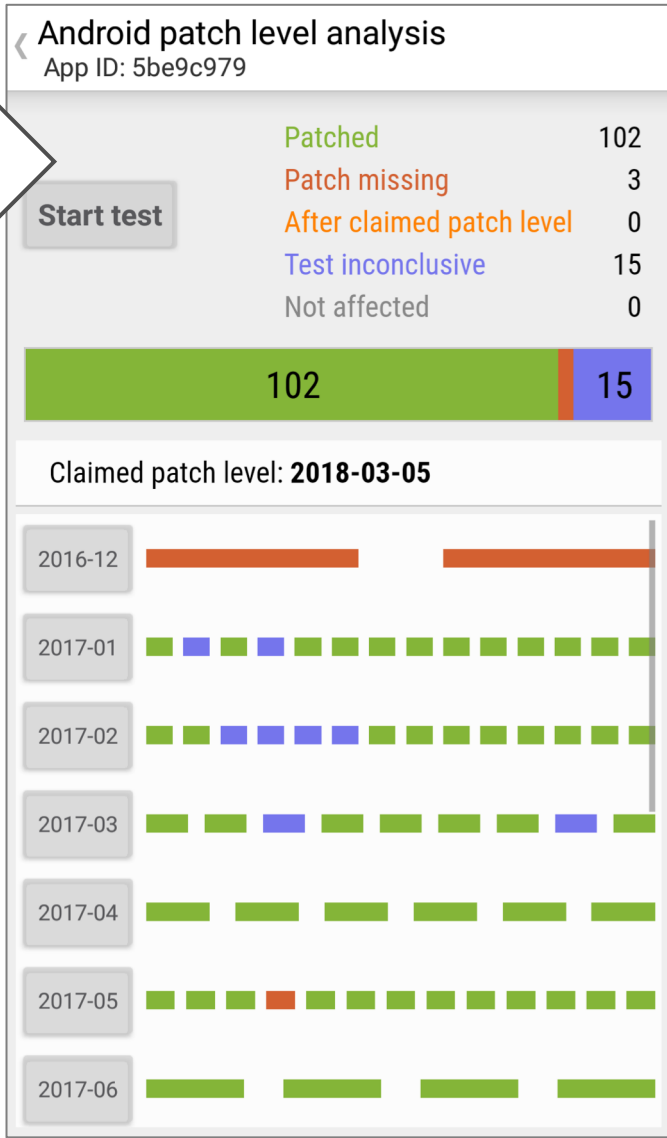
Network Security: Done - You may want to test both 2G and 3G

Provider results in comparison:

	Intercept	Impersonation
Your test result	Higher protection	Lower protection
Your network T-Mobile	Lower protection	Lower protection
Vodafone	Lower protection	Lower protection

	m	w	24 hours	last hour
SMS & SS7 attacks	0	0	0	0
IMSI Catcher	0	0	0	0

Last analysis: 11 Apr 2018 14:10:09



Android patch level analysis
App ID: 5be9c979

Start test

Patched	102
Patch missing	3
After claimed patch level	0
Test inconclusive	15
Not affected	0

102

15

Claimed patch level: 2018-03-05

2016-12	Orange bar
2017-01	Green bar
2017-02	Green bar
2017-03	Green bar
2017-04	Green bar
2017-05	Green bar with red segment
2017-06	Green bar

Take aways

- **Android patching is more complicated and less reliable than a single patch date may suggest**
- **Remote Android exploitation is also more much complicated than commonly thought**
- **You can finally check your own patch level thanks to binary-only analysis, and the app SnoopSnitch**

Many thanks to Ben Schlabs, Stephan Zeisberg, Jonas Schmid, Mark Carney, Luas Euler, and Patrick Lucey!

Questions?

Jakob Lell <jakob@srlabs.de>
Karsten Nohl <nohl@srlabs.de>

References

1. Federal Trade Commision, **Mobile Security Updates: Understanding the Issues**, February 2018
https://www.ftc.gov/system/files/documents/reports/mobile-security-updates-understanding-issues/mobile_security_updates_understanding_the_issues_publication_final.pdf
2. Duo Labs Security Blog, **30% of Android Devices Susceptible to 24 Critical Vulnerabilities**, June 2016
<https://duo.com/decipher/thirty-percent-of-android-devices-susceptible-to-24-critical-vulnerabilities>
3. Google, **Android Security 2017 Year In Review**, March 2018
https://source.android.com/security/reports/Google_Android_Security_2017_Report_Final.pdf