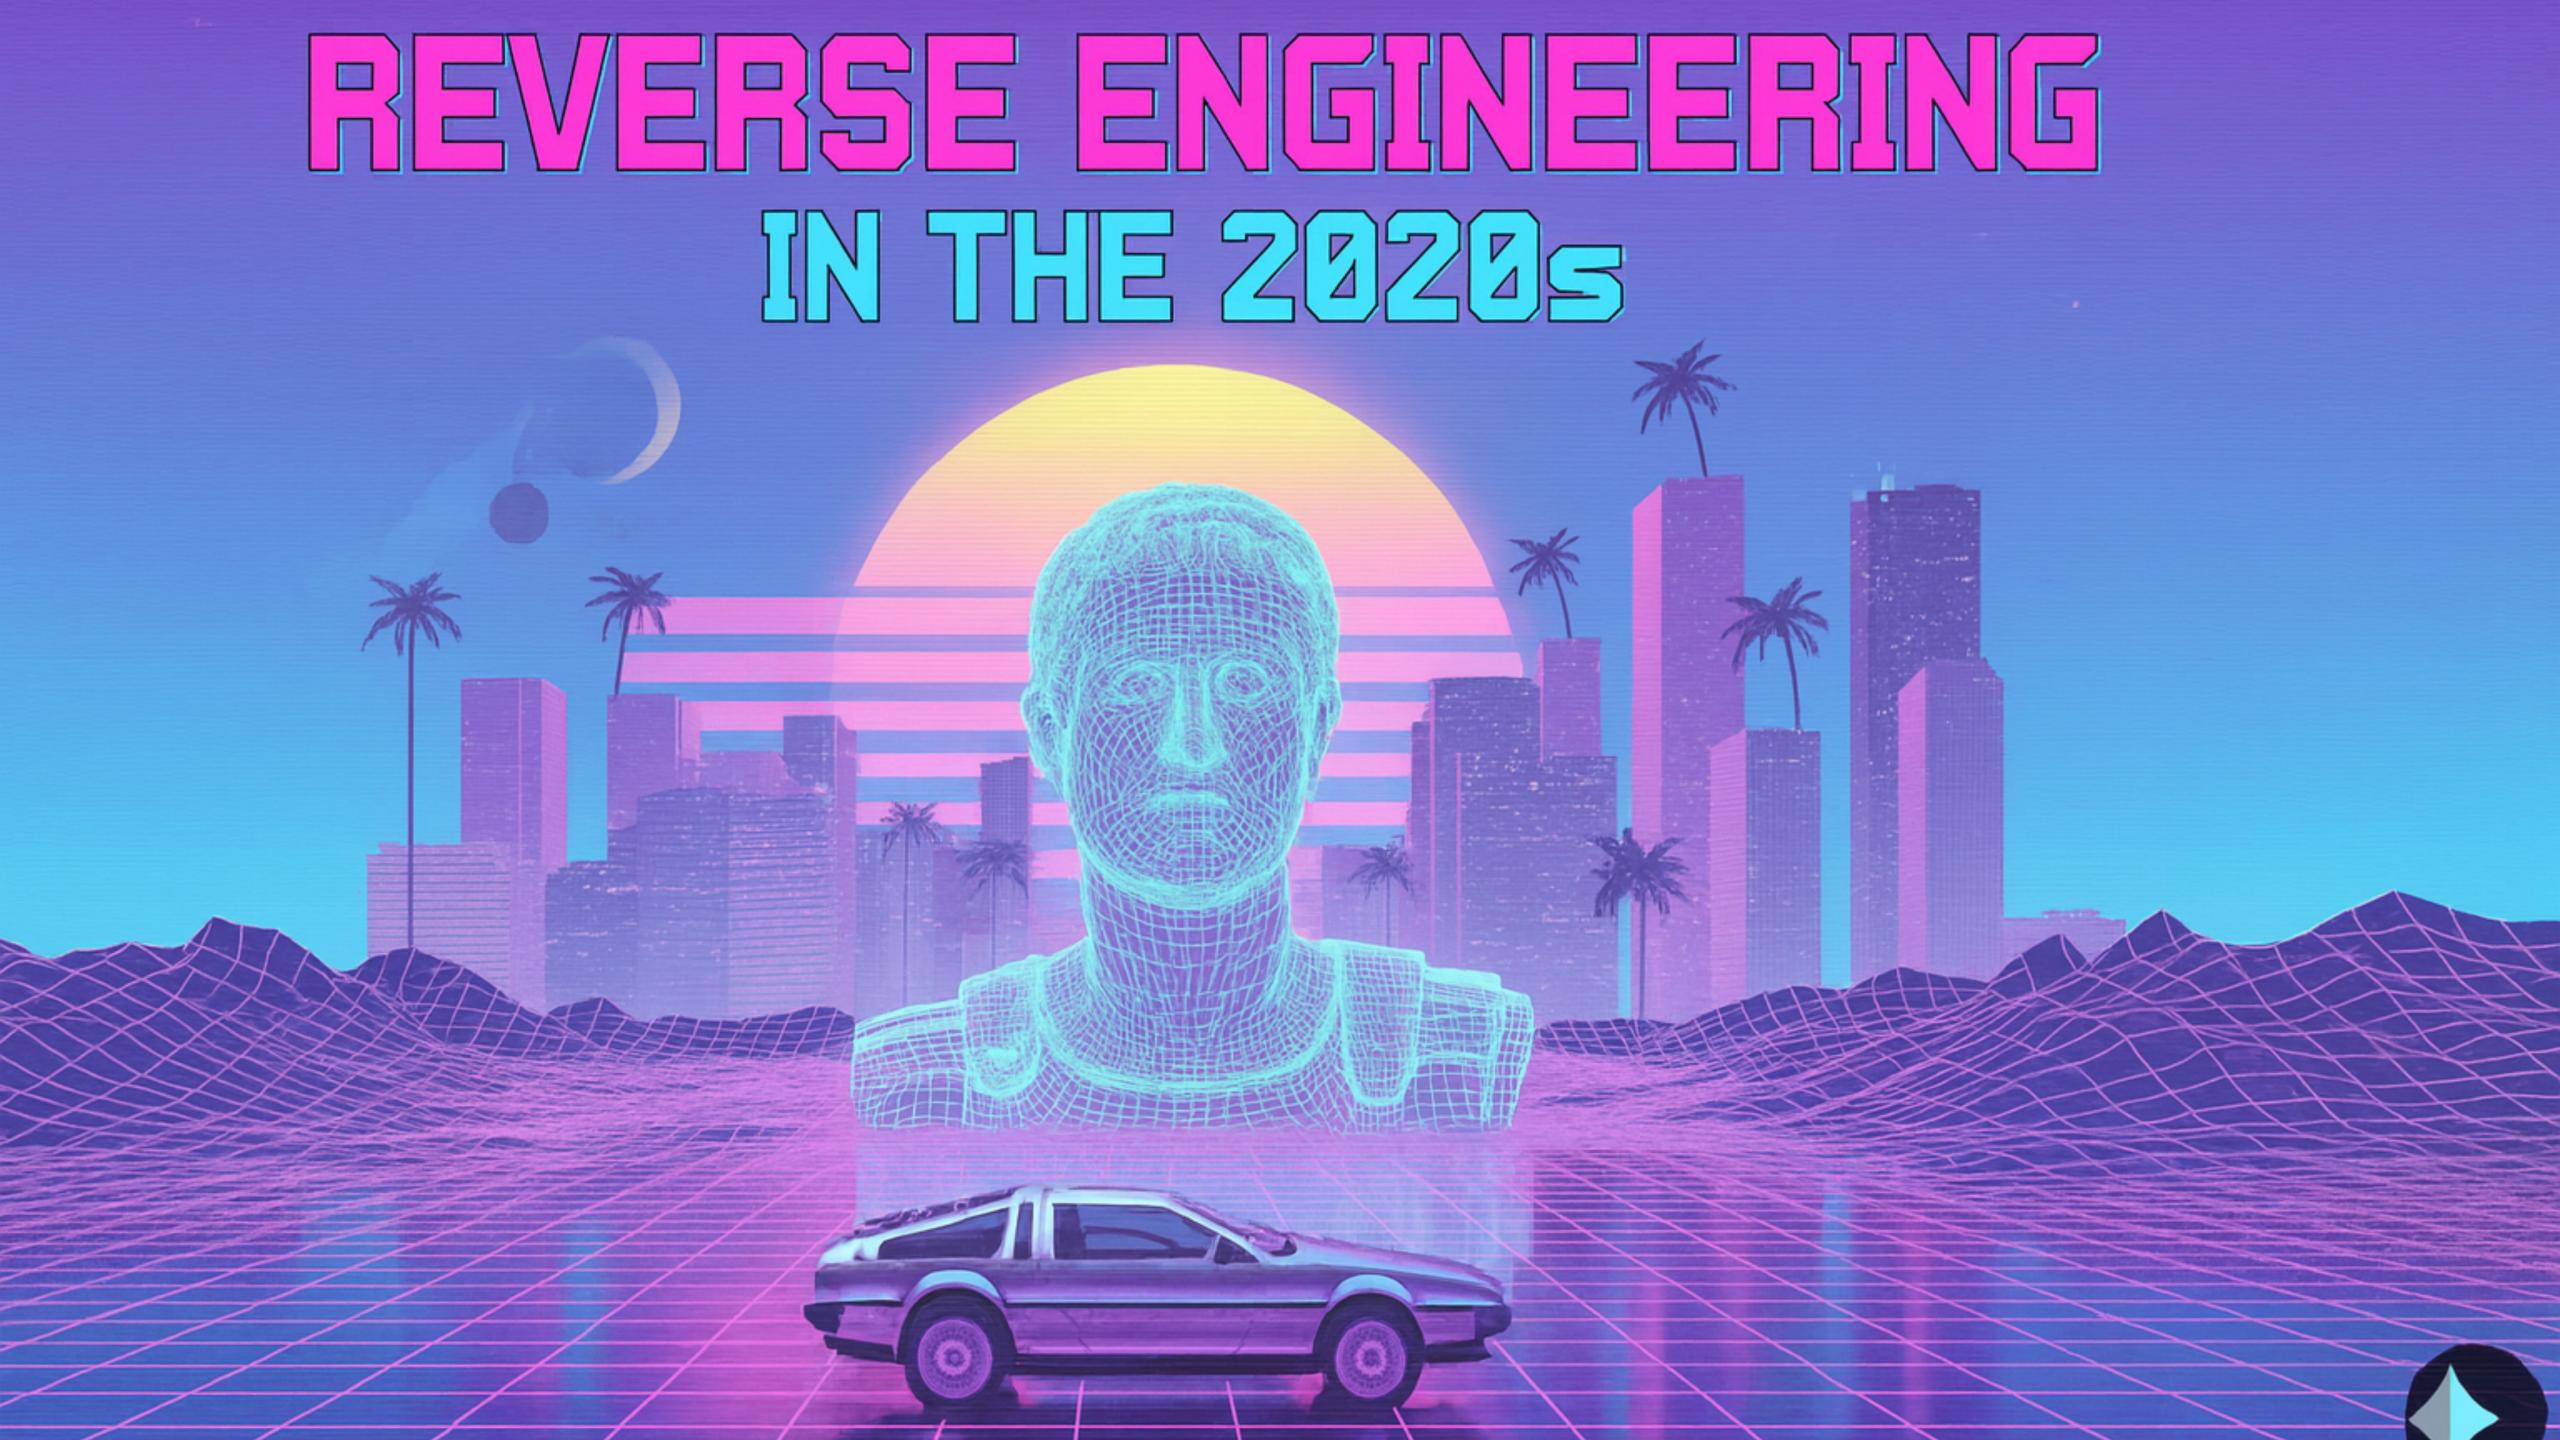


# Reverse Engineering in the 2020s

Where are we, and where should we go?



# Why Reverse Engineering?

I. Motivation

Reverse Engineering (RE) is the process of discovering undocumented internal principles of a piece of code. Why would anyone who is not a criminal want to do that? We

Repeatable Reverse Engineering with PANDA, 2014

- Long ago (the 2000s), the field of reverse engineering had a slightly disreputable reputation
  - Are you trying to steal someone else's intellectual property??
- But there are lots of good reasons to reverse engineer!

**Malware Analysis** 

Interoperability

Software Archaeology

**Vulnerability Discovery** 

Debugging

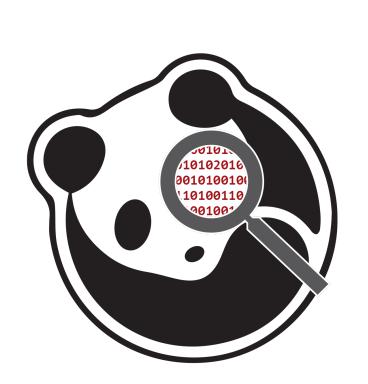
Uncovering Hidden Behavior

### How I Got Here

Started working on reverse engineering to support Windows memory forensics



- Expanded into memory analysis for virtualization security (translation: same thing, but from a hypervisor)
- Built systems for supporting reverse engineering through wholesystem dynamic analyses (PANDA, Malrec)
- Worked on automating reverse engineering for rehosting emulating embedded devices
- Side projects on neural decompilation, fuzzing





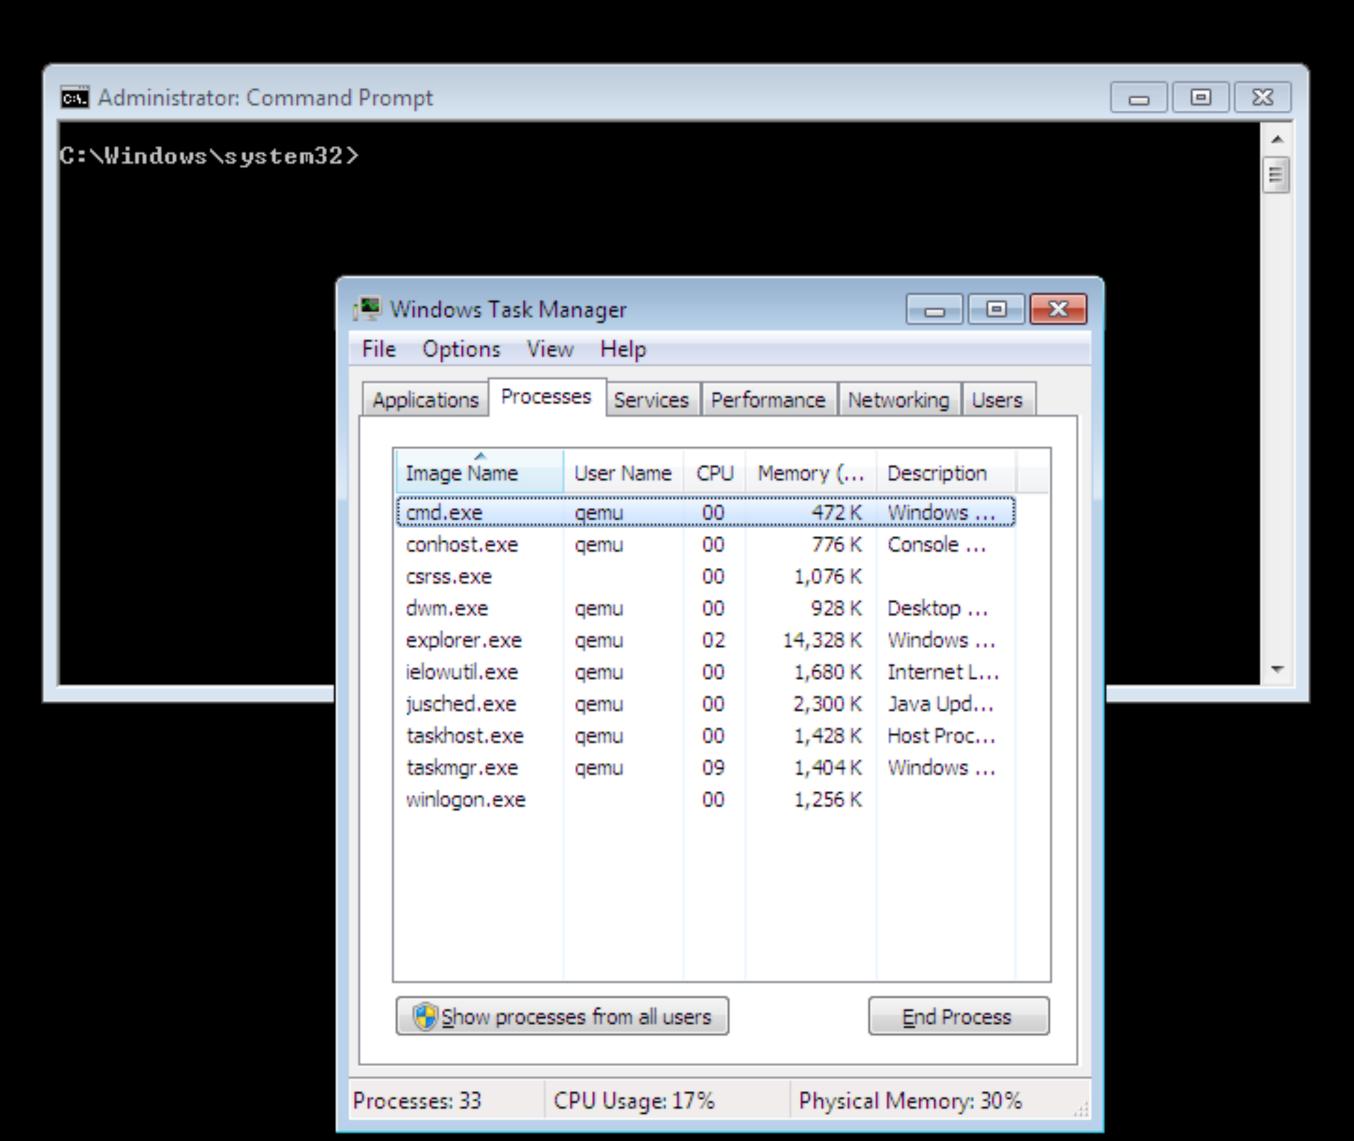
## Reverse Engineering for Memory Forensics

- Basic problem: how do you extract the live state of a system from a memory dump?
- Answer: by reverse engineering the Windows NT kernel, of course!
- You can get lots of cool stuff from this:
  - Lists of running processes, loaded DLLs
  - Information about active network connections
  - Registry data as it appeared in memory

•







Build 7601 This copy of Windows is not genuine







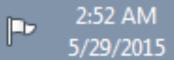






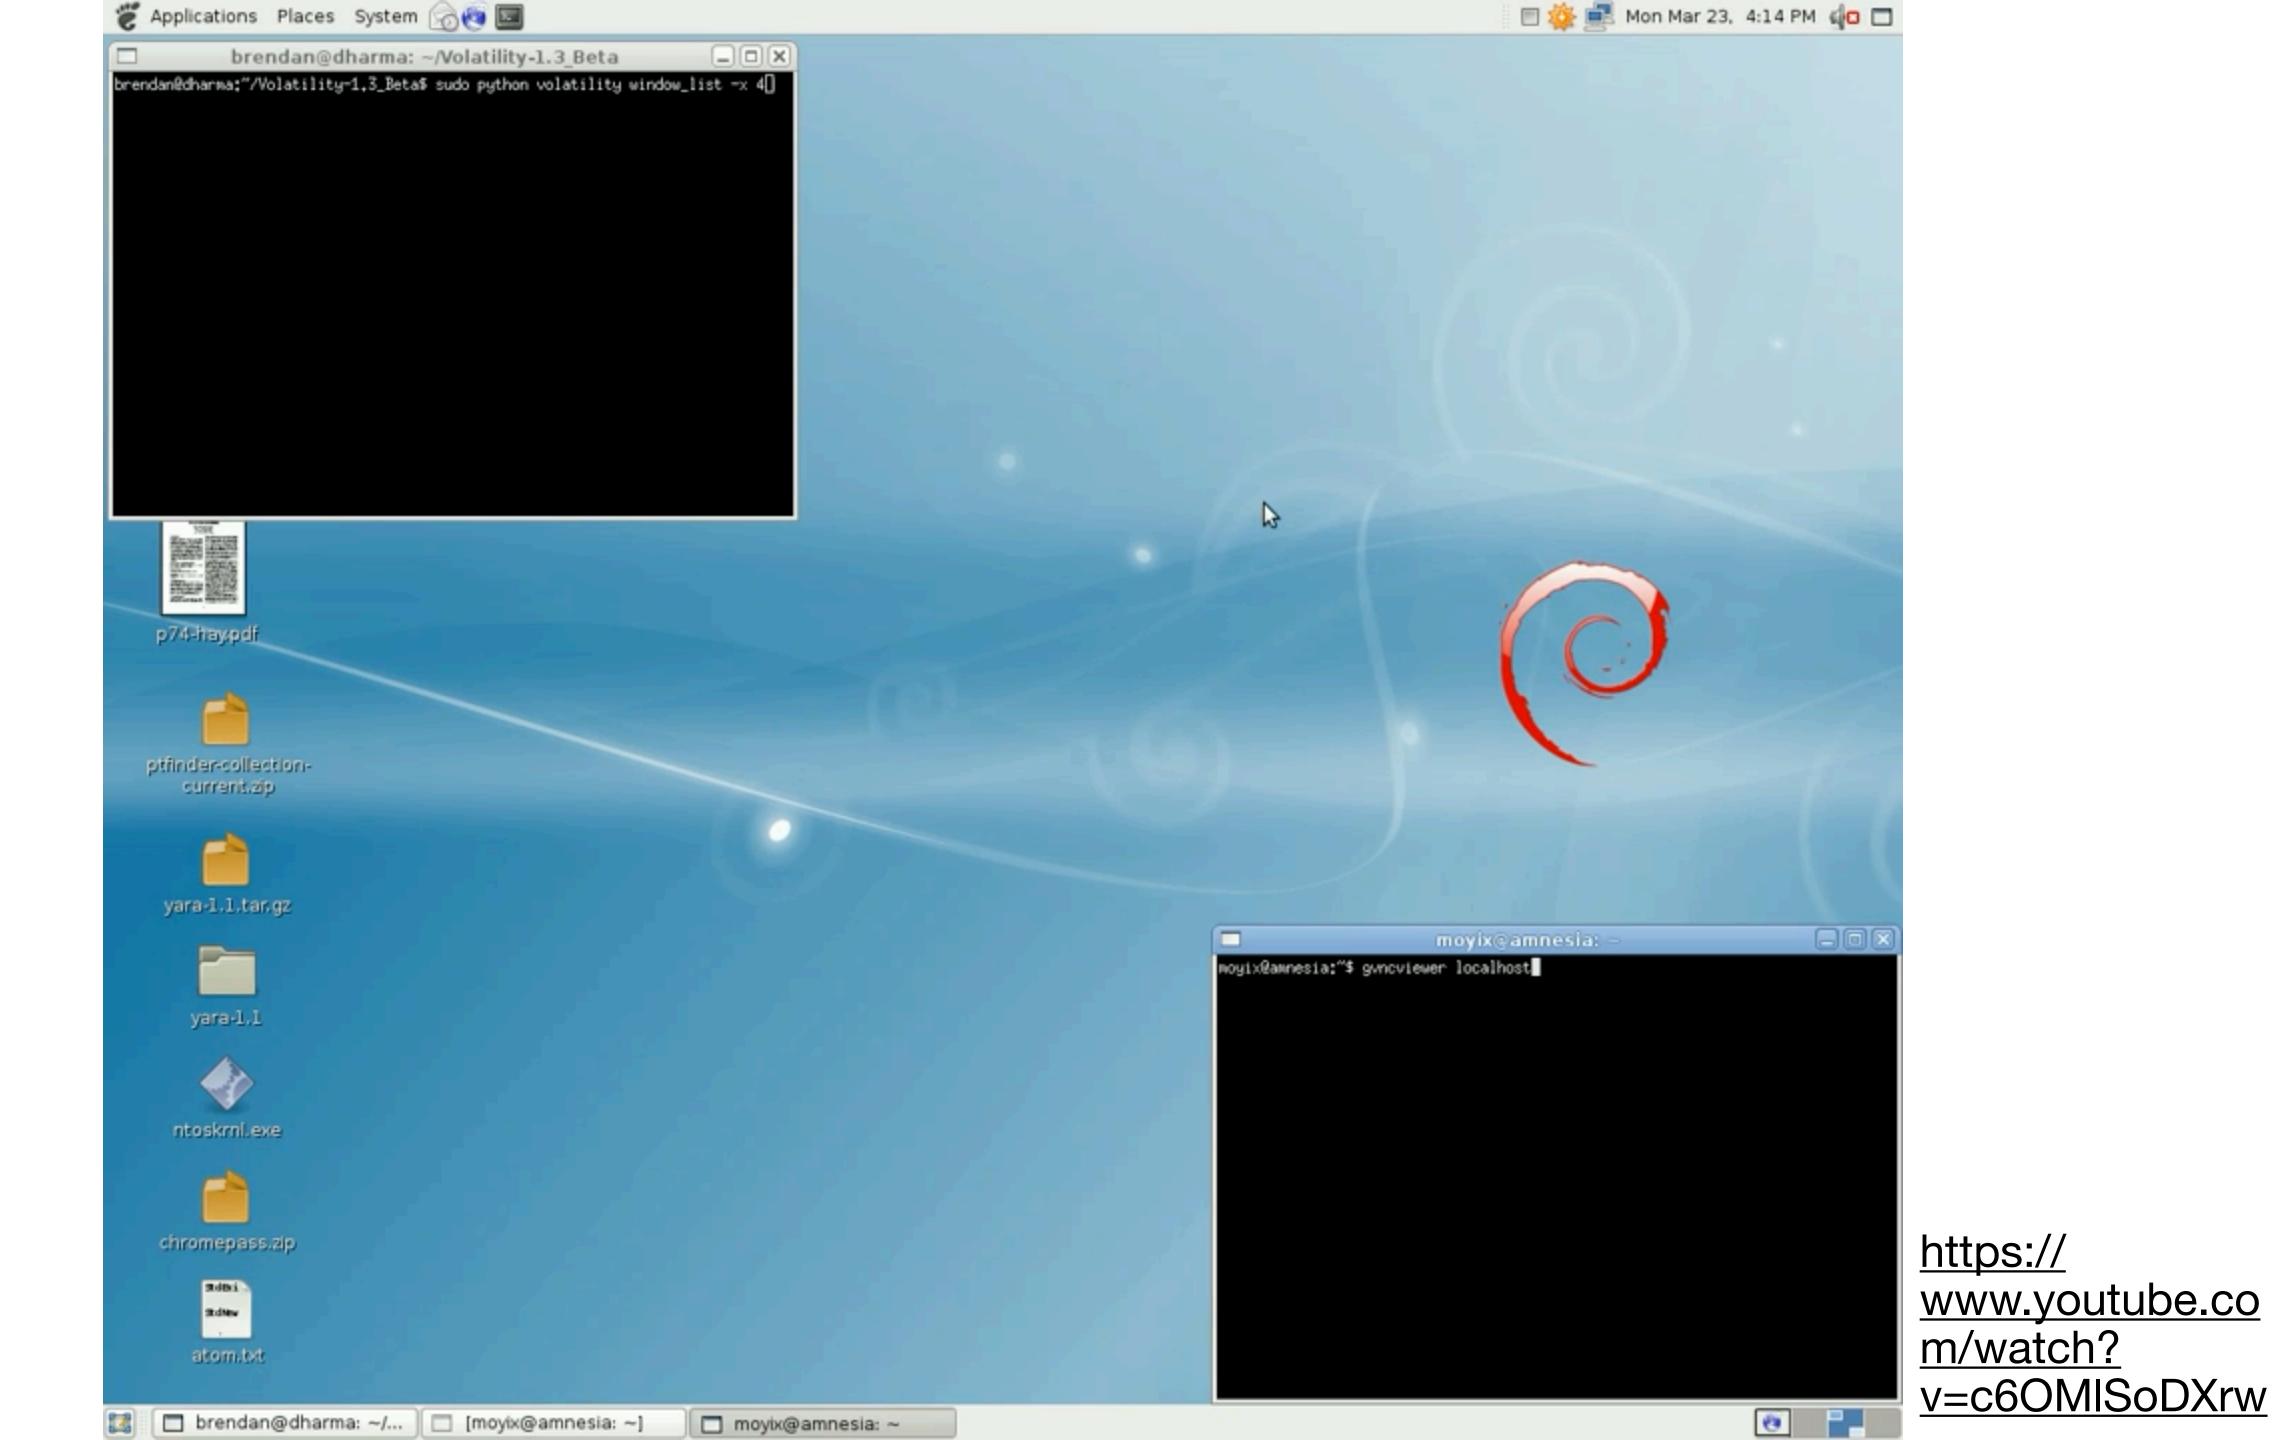






Windows 7

FolderView Administrator: Command Prompt Windows Task Manager Tab1 Processes &Show processes from all users &End Process Running applications start User Promoted2:54 AM



# Whole-System Dynamic Analysis Breaking Spotify DRM

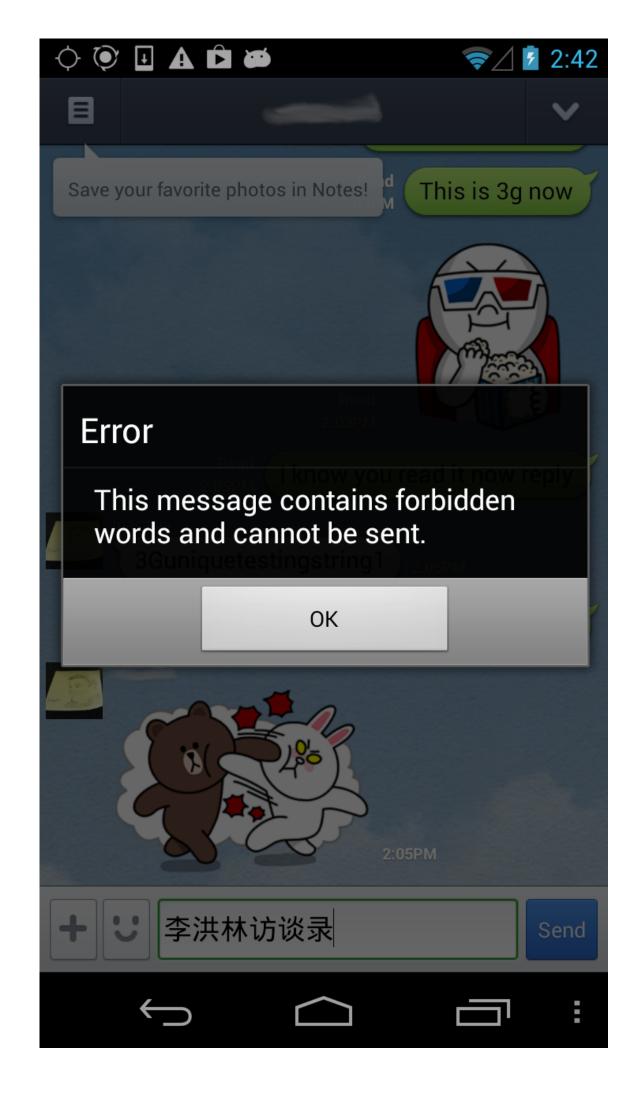
- DRM has a strong signature
  - High entropy, high randomness (χ²) input
  - High entropy, low randomness (χ²) output
- We can look for functions that match this description
- Then just hook their output, and save it!



### Whole-System Dynamic Analysis

#### Censorship Blacklist Extraction

- LINE is a Japanese-made IM app for Android with ~560M users worldwide
- Found by CitizenLab to censor some words for Chinese users
- We wanted to find out which ones
- Very simple strategy: search every memory read & write to find usage of strings likely to be in "bad words" list (法轮 (Falun), 天安门 (Tiananmen))
- Dump everything else passing through that point







# Censorship Blacklist (sample)

	政变	共党	彭博
198964	周斌	共匪	天朝
FLG GCD	祖茔	共	兲朝
GFW	共C档	胡温	屠城
18大	08宪章	江派	屠杀
38军	89事件	江系	团派
八九	艾未未	江	退党
半习习	薄瓜瓜	近平	汪洋
鲍丹乡	薄熙来	九评	瘟神
暴政	曹建明	军警	晓波
柴玲	曾庆红	六四	学潮
赤匪	陈光诚	马凯	学运
	大纪元	民运	余杰

士ノ上井

For translations & context see <a href="https://china-chats.net/">https://china-chats.net/</a>

### Themes in RE Research Over Time

- Over the past 20 years, academic research in reverse engineering has undergone some substantial shifts
  - x86 software running on desktops -> many architectures, desktop/ embedded/mobile
  - Closed tools, hard to reproduce -> open artifacts, good open source tools
  - Reverse engineering to support X -> papers studying RE itself regularly in top conferences

## Then: Closed Tools, Poor Reproducibility

- It used to be very common to release papers with no accompanying code
- This made it extremely difficult to build on prior research
  - Each group needed to reimplement the history of the field before they could make any new progress!
  - Particularly for areas that require significant amounts of system-building, this cost quickly exceeds what anyone is willing to do for a single paper.
- I speculate this is one reason why there was so little progress in decompilation and binary type inference for almost 20 years

### Example: Retypd

#### (Details from Bosamiya et al.'s TRex paper at USENIX Sec 2025)

- Originally presented at PLDI in 2016 by Noonan et al. (GrammaTech)
- In 2021, released open source e
  - ...well, kind of. Actually, reimplemented by a different GrammaTech employee after the original author had left.
  - They found it very difficult!

<sup>1</sup>Indeed for one paper [25], other employees at the same company attempting to reproduce the work stated that "[i]t is a powerful system but difficult to understand" and the "presentation is very dense and subtle" [12].

https://github.com/GrammaTech/retypd/blob/ f8dd231478c3e1722d0d160c3cf99c628a25/reference/type-recovery.rst

### Now: Open Tools and Artifact Evaluation!









TRex: Practical Type Reconstruction for Binary Code

Jay Bosamiya<sup>1\*</sup>, Maverick Woo<sup>2</sup>, and Bryan Parno<sup>2</sup>

<sup>1</sup>Microsoft Research <sup>2</sup>Carnegie Mellon University

# Then: Reverse Engineering for X

- Reverse engineering was most often a supporting discipline
- It was fairly rare to see top conference papers dedicated to RE itself!
- Instead, you would find papers that came up with some cool new technique for binary analysis, and then applied it the topic du jour:
  - Botnets (e.g., automated RE of network protocols)
  - Vulnerability discovery (e.g., CGC & binary fuzzing work)
  - Debloating

This paper is essentially an automated reverse engineering paper disguised as a debloating paper.... and I like it.

Anonymous Reviewer of our IEEE S&P 2022 paper, IRQDebloat

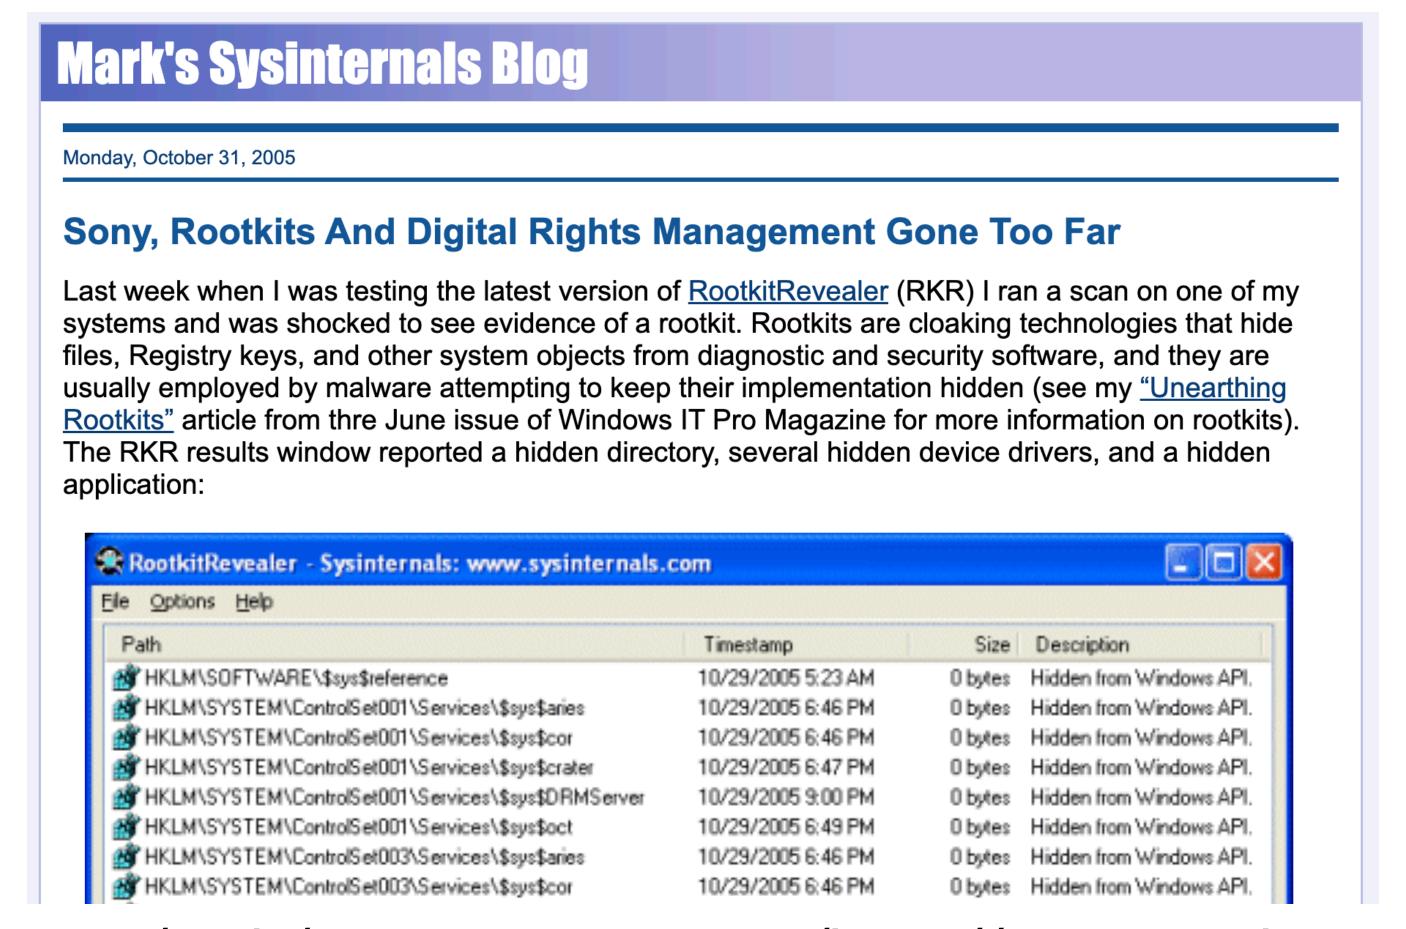
# Now: Reverse Engineering on Its Own

- Increased interest in:
  - Decompilation: <a href="https://mahaloz.re/dec-progress-2024">https://mahaloz.re/dec-progress-2024</a>
  - Type reconstruction
  - Reliable disassembly (and reassembly)
  - Fuzzing of binary software
  - Automated reverse engineering of embedded devices (rehosting)

### Fulfilling the Potential of RE?

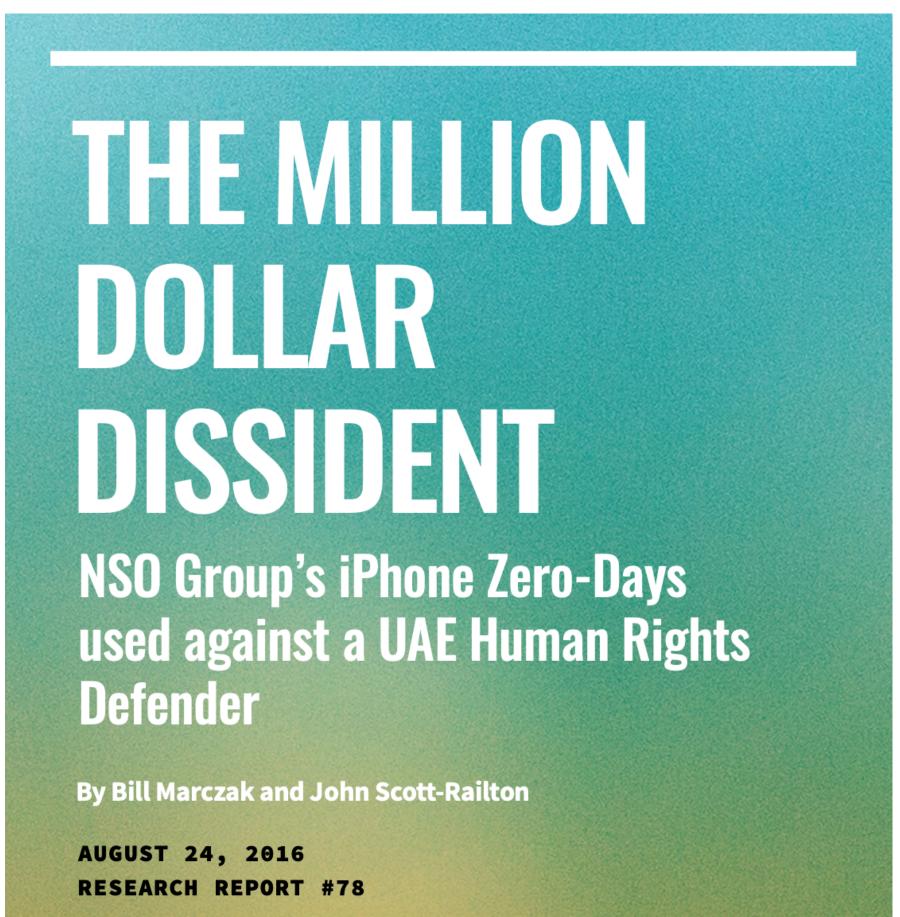
- Reverse engineering has the potential to help return control of software back to its users
- This is important since software increasingly mediates so much of what we do every day: how we talk to other people, the devices in our homes, even the code that runs on cars and trains
- We have seen many examples of how reverse engineering can be used as a force for good, uncovering bad behavior in software systems

Sony Rootkit (2005)



https://web.archive.org/web/20051102053346/http://www.sysinternals.com/blog/2005/10/sony-rootkits-and-digital-rights.html

CitizenLab Reverse Engineering of Malware Targeting Activists



https://citizenlab.ca/2016/08/million-dollar-dissident-iphone-zero-day-nso-group-uae/



Figure 3: SMS text messages received by Mansoor (English: "New secrets about torture of Emiratis in state prisons"). The sender's phone numbers are spoofed.

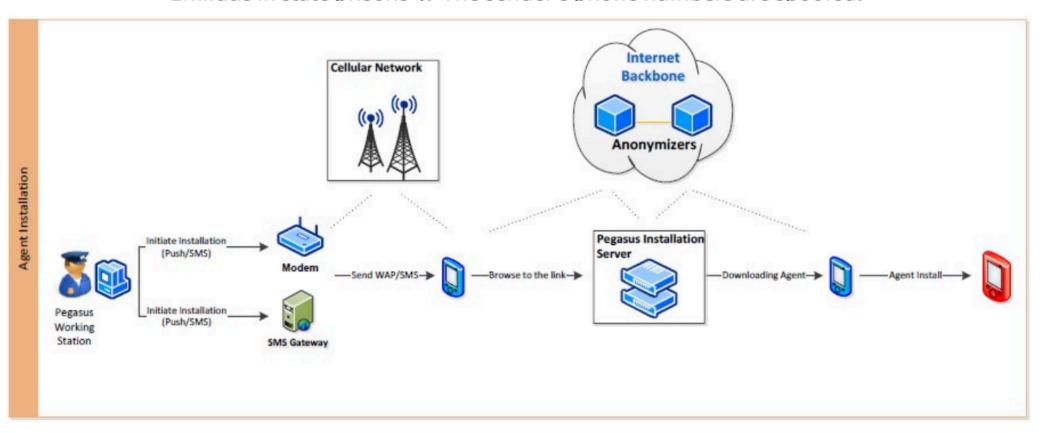
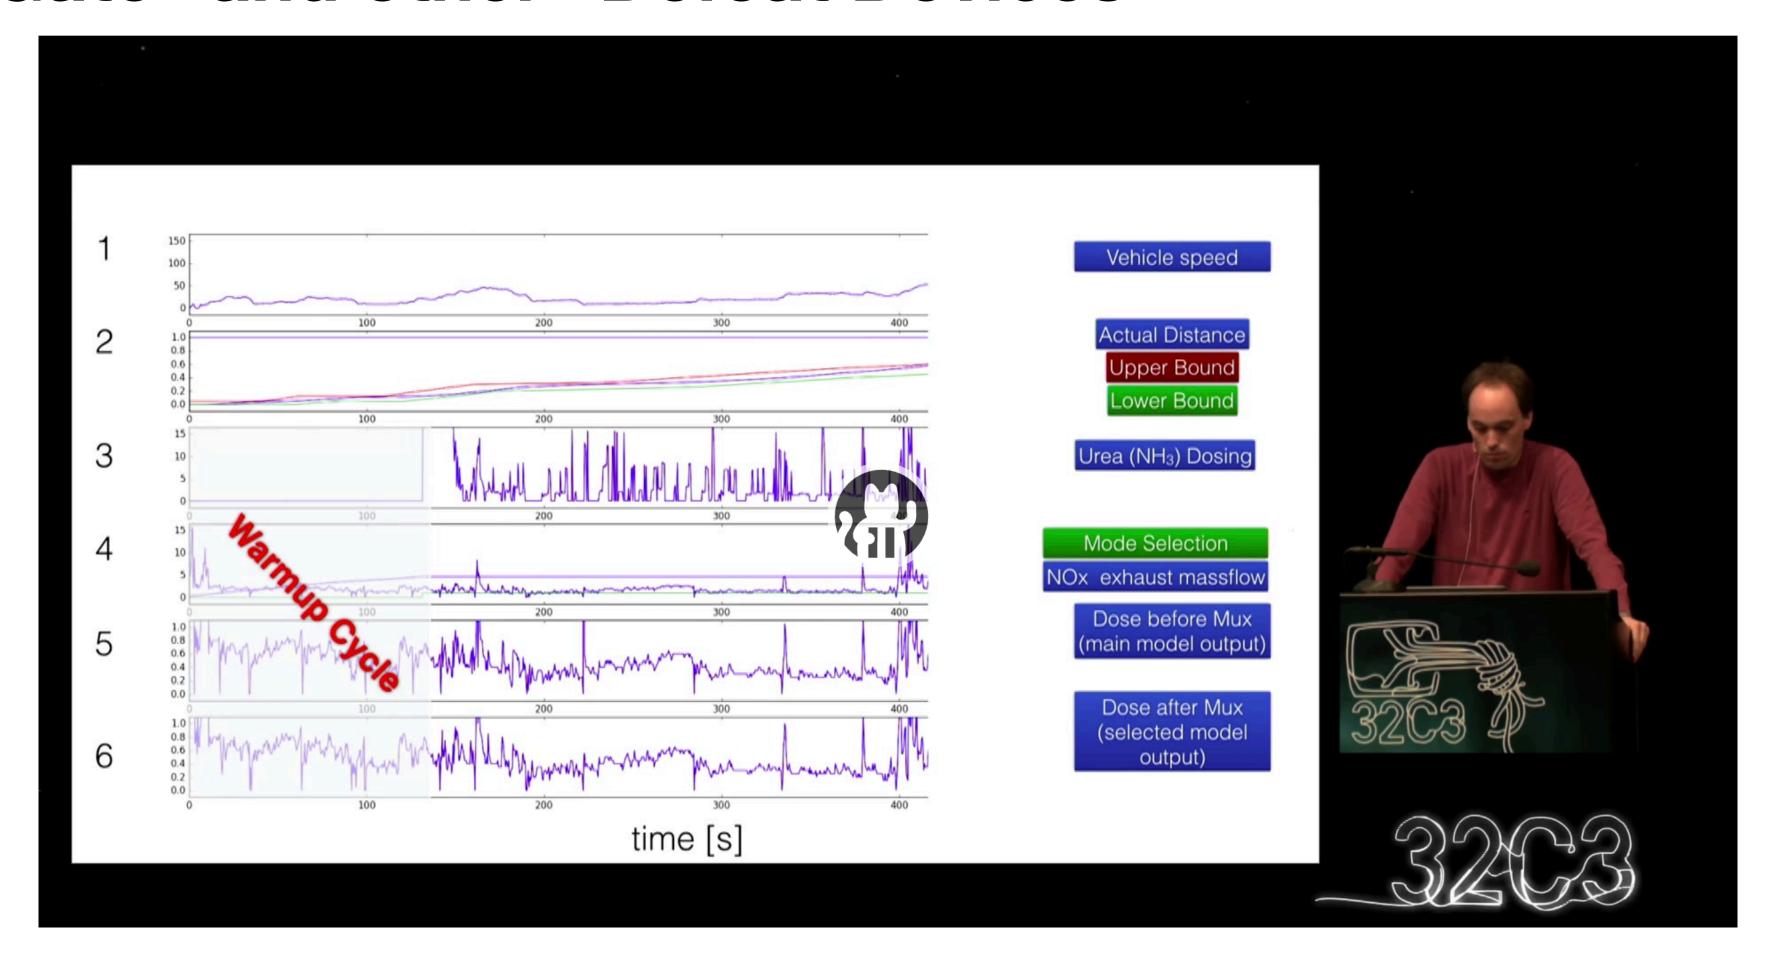


Figure 6: Diagram from purported NSO Group Pegasus documentation showing the sequence through which the spyware ("Agent") is installed on a target's mobile device. Source: Hacking Team Emails.

"DieselGate" and other "Defeat Devices"



https://media.ccc.de/v/32c3-7331-the exhaust emissions scandal dieselgate

#### "DieselGate" and other "Defeat Devices"

## How They Did It: An Analysis of Emission Defeat Devices in Modern Automobiles

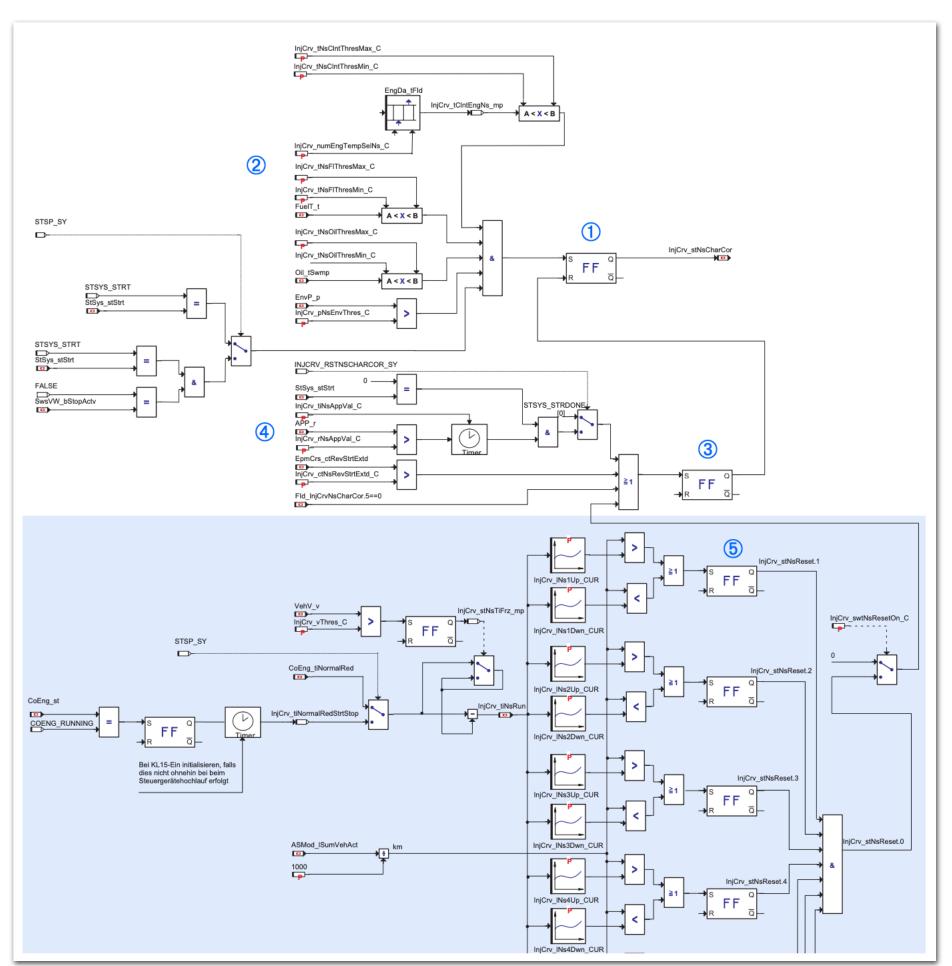
Moritz Contag\*, Guo Li<sup>†</sup>, Andre Pawlowski\*, Felix Domke<sup>‡</sup>, Kirill Levchenko<sup>†</sup>, Thorsten Holz\*, and Stefan Savage<sup>†</sup>

\* Ruhr-Universität Bochum, Germany, {moritz.contag, andre.pawlowski, thorsten.holz}@rub.de

† University of California, San Diego, {gul027, klevchen, savage}@cs.ucsd.edu

† tmbinc@elitedvb.net

In this paper, we present our analysis of two families of software defeat devices for diesel engines: one used by the Volkswagen Group to pass emissions tests in the US and Europe, and a second that we have found in Fiat Chrysler Automobiles. To carry out this analysis, we developed new static analysis firmware forensics techniques necessary to automatically identify known defeat devices and confirm their function. We tested about 900 firmware images and were able to detect a potential defeat device in more than 400 firmware images spanning eight years. We describe the precise conditions used by the firmware to detect a test cycle and how it affects engine behavior. This work frames the technical challenges faced by regulators going forward and highlights the important research agenda in providing focused software assurance in the presence of adversarial manufacturers.



https://cseweb.ucsd.edu/~klevchen/diesel-sp17.pdf

#### DragonSec RE of Anti-Repair Code in Newag Train Firmware



https://zaufanatrzeciastrona.pl/post/o-trzech-takich-co-zhakowali-prawdziwy-pociag-a-nawet-30-pociagow/

### But: RE is Not Accessible

- All of the (extremely cool) analyses just discussed required enormous amounts of work by dedicated specialists
- There is too much software in the world, and too few reverse engineering experts, for this to scale
- Even most "automated" analyses require significant expertise to use and may not scale to real software





Speculating wildly, it feels like we may be getting close to something I've dreamed about since grad school – making the software running on your computer inspectable and understandable to the end user.



Entirely anecdotal, but ChatGPT+GPT-4 was \*extremely\* helpful recently for figuring out why a closed-source program on Mac was crashing. I just attached Xcode to it and fed in chunks of assembly, asking what the original ObjC/Swift code might have been, until I hit root cause.

7:51 PM · Mar 31, 2023 · **8,977** Views



Brendan Dolan-Gavitt 📀 @moyix · Mar 31, 2023



Imagine an interface that let you point to something and say "what does this button do?" and it would pull out the relevant ASM, use GPT-4 (+trad reverse engineering tools?) to summarize its functionality, and produce an explanation linked to the original assembly.

0

17

**O** 19

ıl<sub>ı</sub>ı 1.6K

 $\supset$ 

Λ,



Brendan Dolan-Gavitt 📀 @moyix · Mar 31, 2023



You could maybe also combine this with some kind of whole-system record/replay (with a smallish event buffer) so that you could ask questions events that happened on your system, like "Why did that window just pop up?"

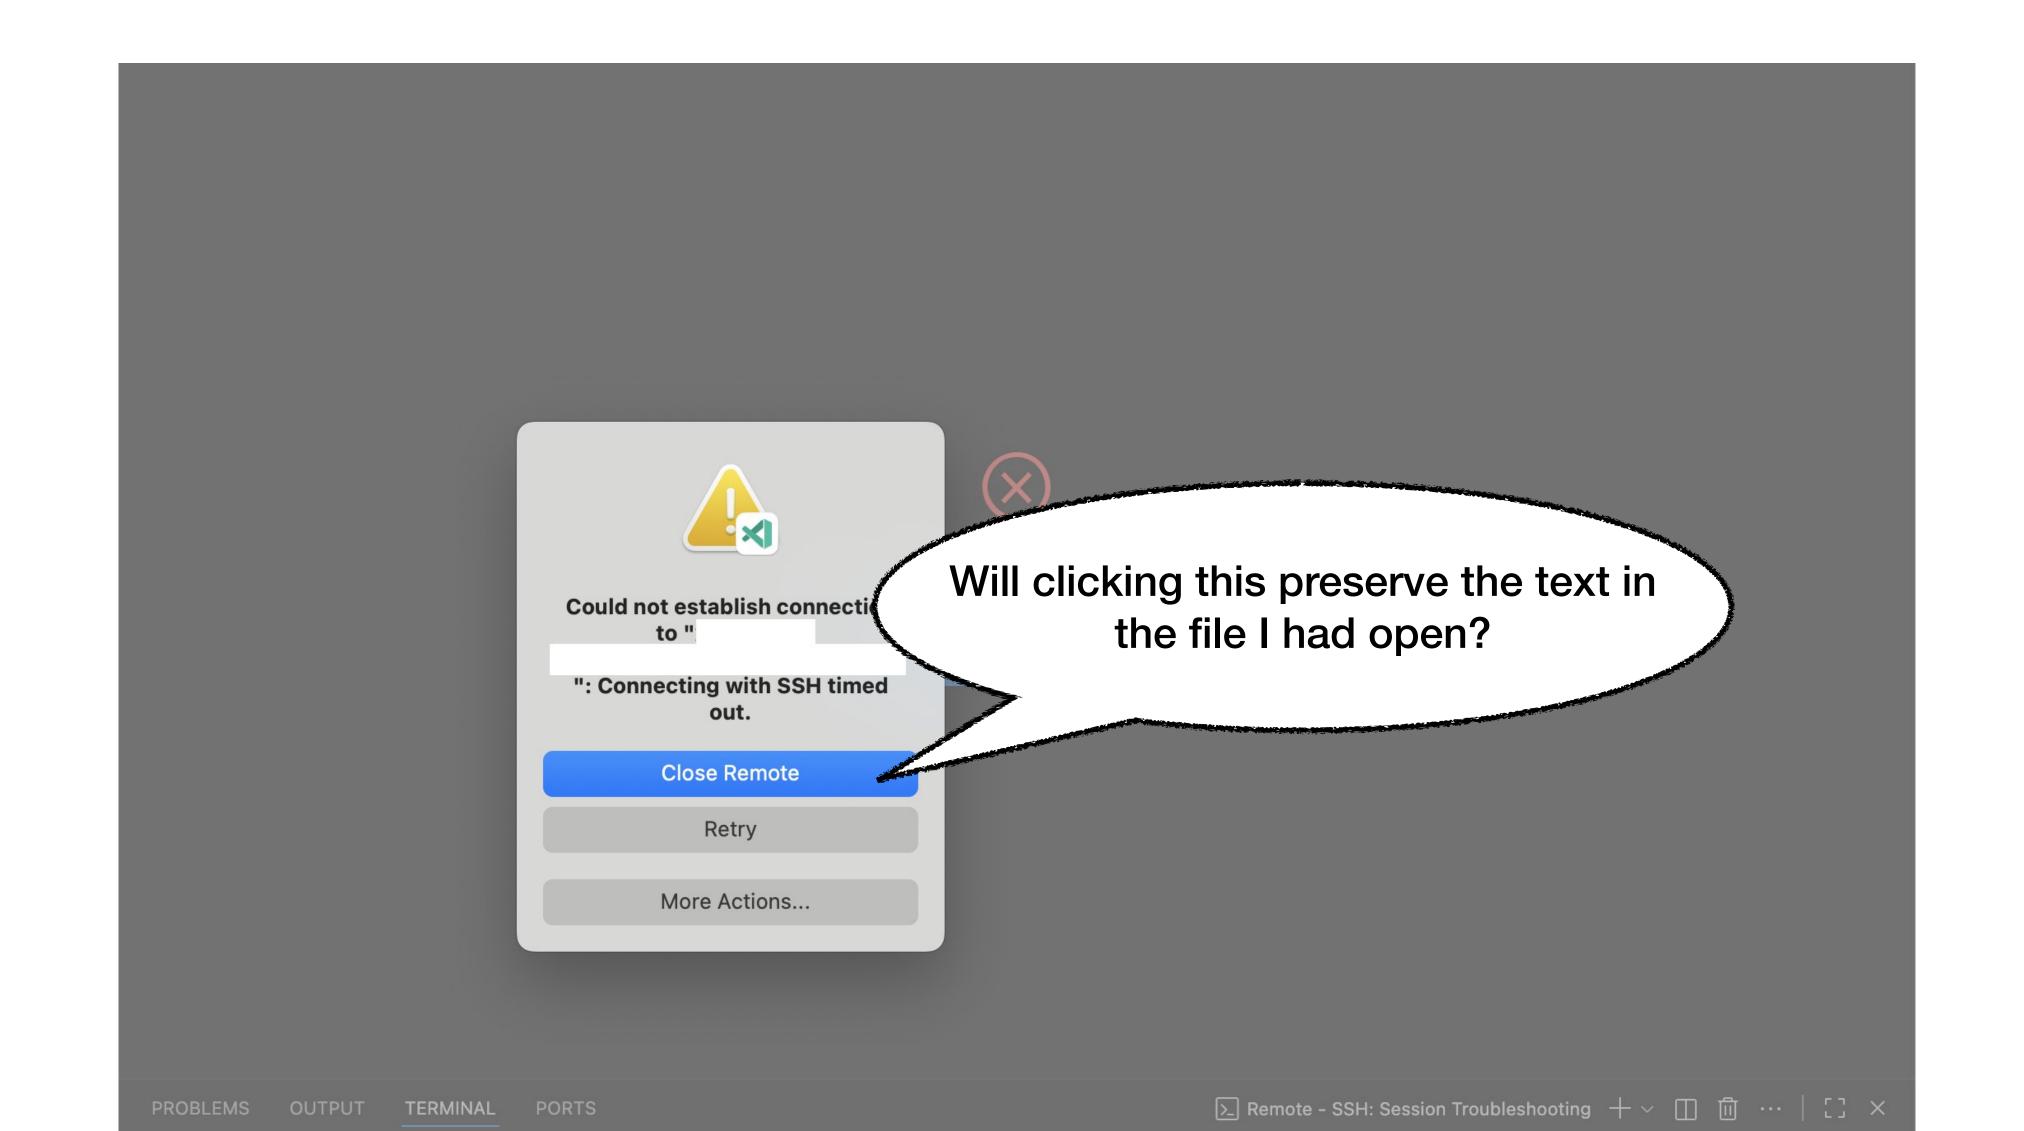
**Q** 2

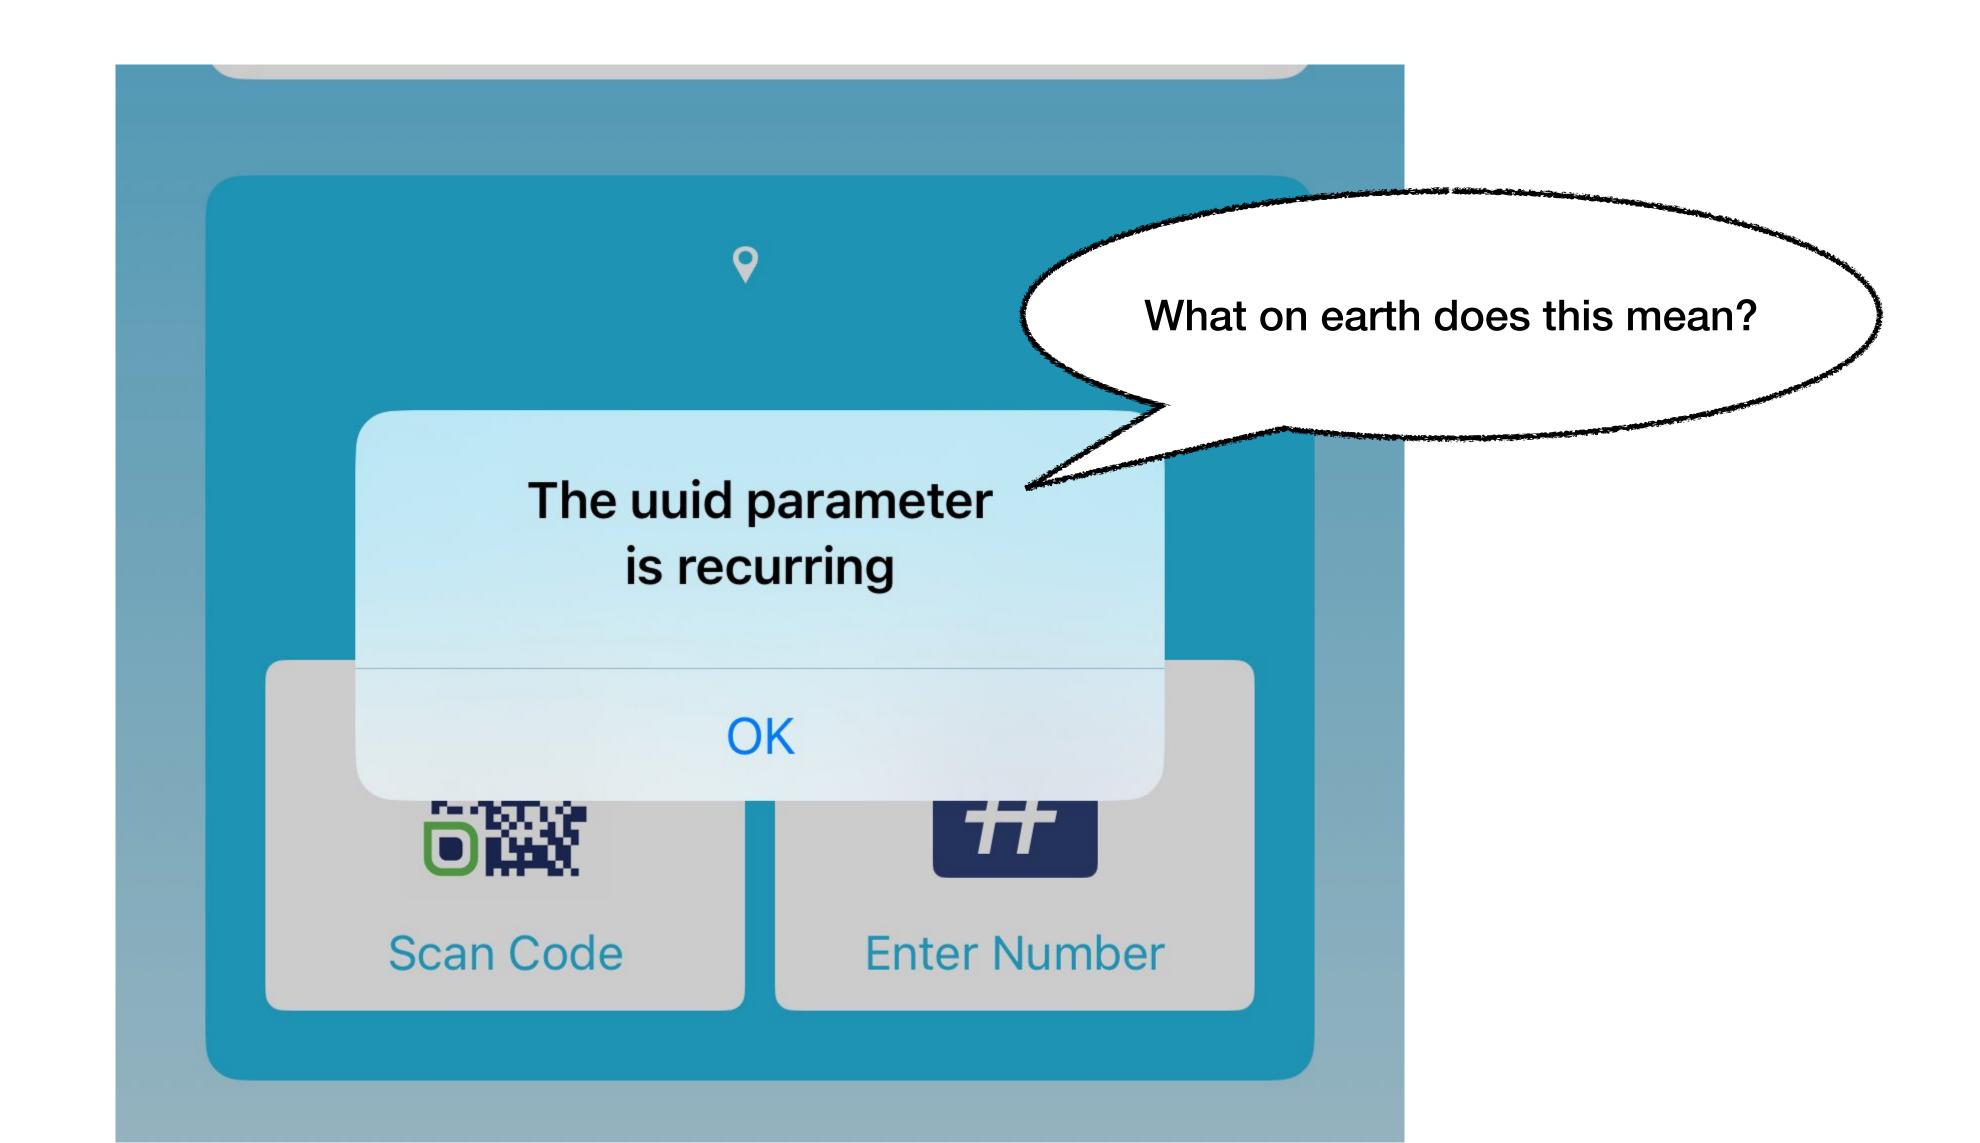
17

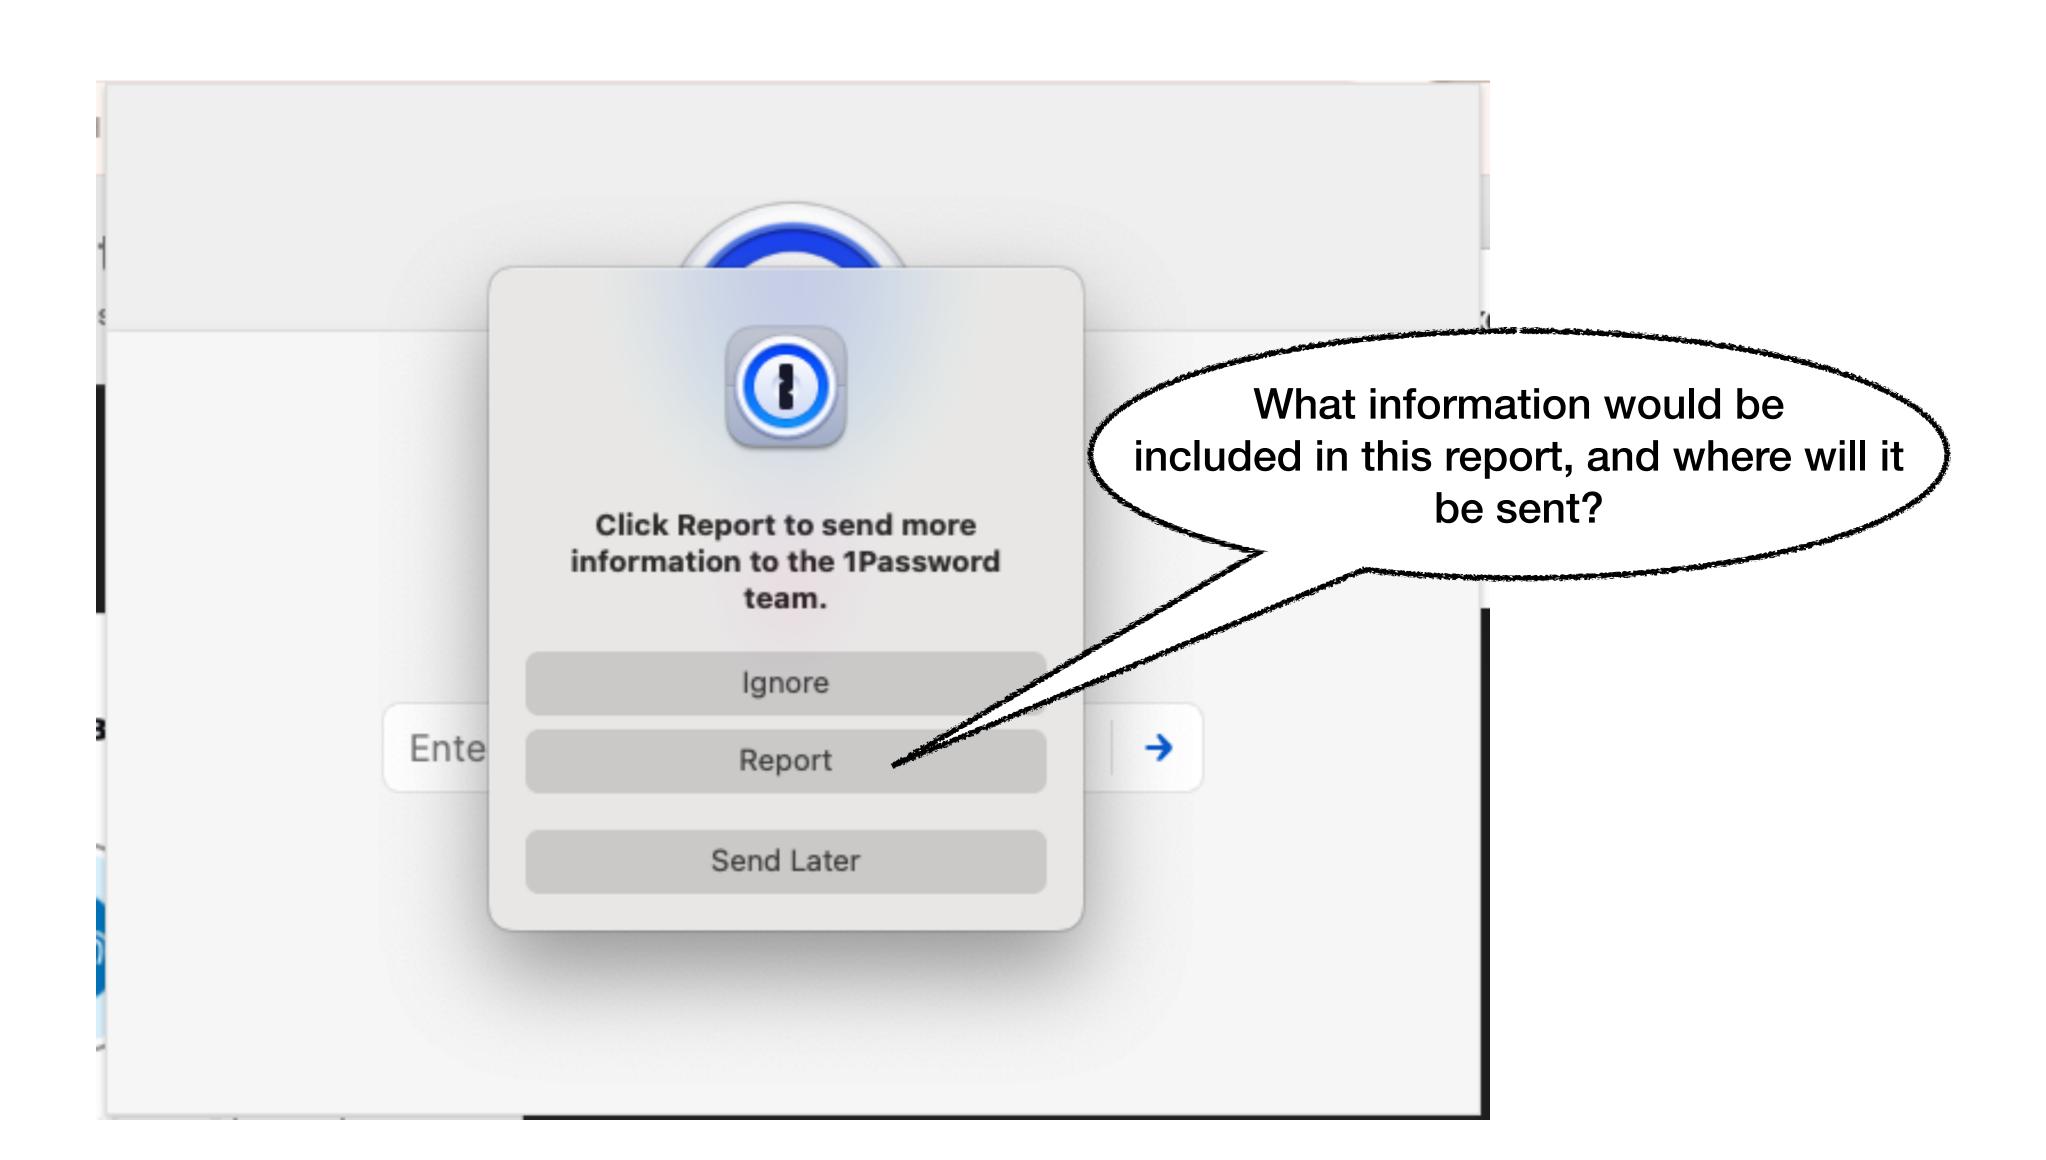
**)** 15

ılıı 1.4K

] ,







## Early Signs of Life?

Full transcript: https://chat.openai.com/share/b41bcd92-80f8-4aa4-9086-e8aede516ace

 Anecdote: I was doing my taxes with TurboTax (closed source) on OS X, but it kept crashing at a particular point



- Using ChatGPT4, I was able to:
  - Get instructions for attaching the XCode debugger and getting a backtrace
  - Feed the backtrace to GPT4 to identify the problematic code
  - Decompile functions along the way to Objective C / Swift
  - Identify and fix the problem (a missing JSON file in the installation)
- (I don't have much experience with OS X or Swift/Objective C reverse engineering)

### Crash Source Identification



The crash occurs on the main thread (thread #1) with the stop reason: `Fatal error:

Unexpectedly found nil while implicitly unwrapping an Optional value`. The issue seems to be related to the unwrapping of an optional value without proper checking for `nil`.

The relevant stack trace line is:

`frame #6: 0×000000107814d1d TurboTax 2020`static

TurboTax\_2020.FDPContingencyHandler.readAndReturnExperimentInConfig() -> Swift.Dictionary<Swift.String, Swift.AnyObject> + 637`

This line is within the

`TurboTax\_2020.FDPContingencyHandler.readAndReturnExperimentInConfig()`

function, which returns a dictionary. It appears that the crash is due to an implicitly unwrapped optional value being 'nil' when it was expected to have a value. To resolve this issue, it's necessary to find out why the optional value is 'nil' and ensure proper handling of optional values within the code. This can be done by using optional binding (if let or guard let) or providing a default value with the nil-coalescing operator (??).

# Swift Decompilation



It's quite challenging to reconstruct the exact original Swift code from assembly, but I can provide you with a rough outline of what the code might have looked like. The details and syntax might be different from the actual code. Here's an attempt to reconstruct the code:

```
syntax might be different from the actual code. Here's an attempt to reconstruct the code:
                                                                     Copy code
 swift
  class TurboTax_2020 {
     // ...
     static func FDPContingencyHandler() -> [String: AnyObject] {
          var dictionary: [String: AnyObject] = [:]
          let experimentServiceController = ExperimentServiceController()
          let experimentSettings = experimentServiceController.loadExperimentS
         if let settings = experimentSettings {
              // Code to process experimentSettings and populate the dictionar
              // ...
         return dictionary
```

# Making Binary Systems (More) Transparent

- I don't actually think LLMs are ready to do all of this on their own
- More reliable, deductive systems are still needed to provide the meat of the analyses
  - Lightweight record-replay to keep track of things that are happening
  - Control and data flow analyses to understand which components are involved and identify the code that needs to be examined
  - Semantics-preserving decompilation
- But LLMs may provide helpful glue at the input and output portions:
  - Turning user queries into a sequence of calls that carry out automated analyses on binary code
  - Digesting the results of reliable analyses into an understandable (but accurate!) answer

## Challenges (Research-Oriented)

- Real systems are a complicated mess of different technologies
  - Native libraries, JS code, bytecode, IPC, calls to remote servers
  - Even for binary code the "style" of the object code can vary significantly depending on source language, application domain, etc.
- Performant whole-system record/replay basically doesn't exist
  - Though with some hardware extensions it seems feasible
- How can we make sure the LLM bits faithfully translate what the deterministic/symbolic components provide?
- Tons of interesting UX and agent design challenges in putting all the pieces together and making them work on real software

# Challenges (non-technical)

- Who will pay for this? Not much of a business model.
- Requires significant device privilege to collect and inspect relevant data
  - Trend has been toward taking this away from end users!
- If it becomes popular, we can expect anti-RE measures to become more common in "regular" binary software
  - But maybe not ubiquitous? The web is currently very "inspectable" by default, but obfuscation is still not the norm

### Conclusions

- Reverse engineering has come a long way since I started my PhD in 2008
- We've gone from a culture of unreproducible, specialized one-offs to powerful tools that work on real software and can be built upon
- I am very excited and optimistic about the potential to finally take automated reverse engineering out of the lab and put it in the hands of people who can benefit from it most
  - A bit sad that I probably won't be the one doing most of this work!
  - But reassured to see a workshop like this carrying on the torch of excellent research in binary analysis