





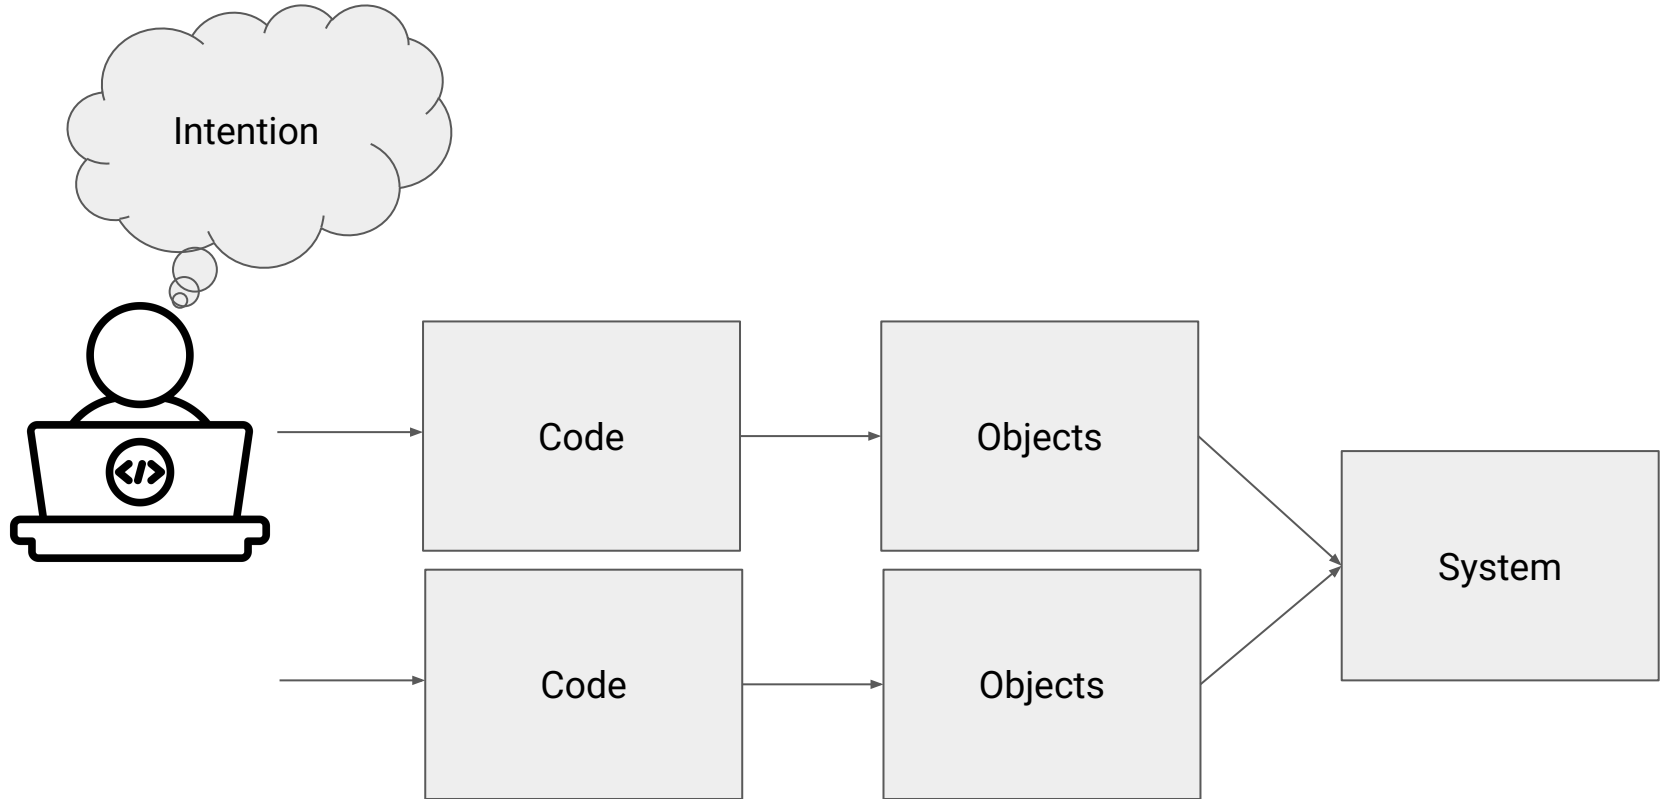
VULCAN: Vulnerable Logic Discovered with Automated Intent Analysis

Arizona State University
INGOTS Hackathon
5/7/26

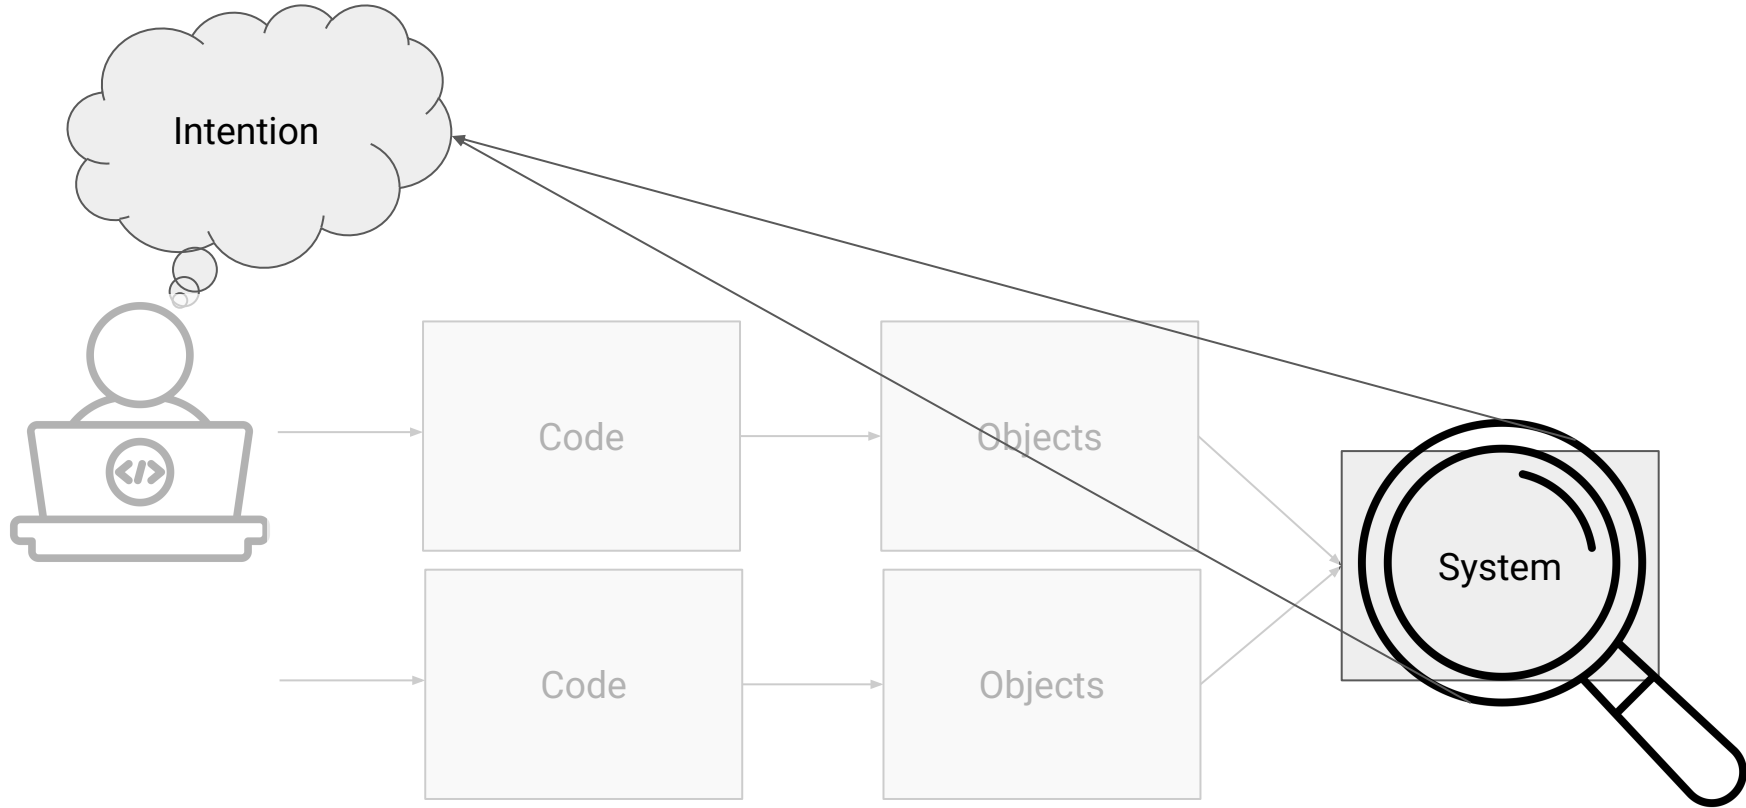
Team Introduction



Logic Bugs



Detecting Logic Bugs



Active Research Area.. For Web Applications

Toward Automated Detection of Logic Vulnerabilities in Web Applications

Viktorija Felmetzger Ludovico Cavedon Christopher Kru
[rusvika,cavedon,chris.vigna]@cs.ucsb.edu
Computer Security Group
Department of Computer Science
University of California, Santa Barbara

Abstract

Web applications are the most common way to make services and data available on the Internet. Unfortunately, with the increase in the number and complexity of these applications, there has also been an increase in the number and complexity of vulnerabilities. Current techniques to identify security problems in web applications have mostly focused on input validation flaws, such as cross-site scripting and SQL injection, with much less attention devoted to application logic vulnerabilities.

Application logic vulnerabilities are an important class of defects that are the result of faulty application logic. These vulnerabilities are specific to the functionality of particular web applications, and, thus, they are extremely difficult to characterize and identify. In this paper, we propose a first step toward the automated detection of application logic vulnerabilities. To this end, we first use dynamic analysis and observe the normal operation of a web application to infer a simple set of behavioral specifications. Then, leveraging the knowledge about the typical execution paradigm of web applications, we filter the learned specifications to reduce false positives, and we use model checking over symbolic input to identify program paths that are likely to violate these specifications under specific conditions, indicating the presence of a certain type of web application logic flaws. We developed a tool, called *Water*, based on our ideas, and we applied it to a number of web applications, finding previously-unknown logic vulnerabilities.

financial constraints. Application vulnerabilities flected in the Symantec Report, which was published states that, in 2007 for 63% of the total number of vulnerabilities. Current techniques to identify security problems in web applications have mostly focused on input validation flaws, such as cross-site scripting and SQL injection, with much less attention devoted to application logic vulnerabilities.

Most recent research applications has focused on characterizing vulnerabilities in web applications. This research uses external input as an out first checking or examples of input values (XSS) [20] and SQL injection. With XSS, an application not sufficiently check malicious JavaScript code injected into the application. This tool found 3,944 EAR instances in 18,127 open-source applications. Finally, we describe an approach to prevent EARs in web frameworks.

One reason for the vulnerabilities is that it is a general specification: characteristics of these vulnerabilities that represent security concerns. Then, we use model checking over symbolic input to identify program paths that are likely to violate these specifications under specific conditions, indicating the presence of a certain type of web application logic flaws. We developed a tool, called *Water*, based on our ideas, and we applied it to a number of web applications, finding previously-unknown logic vulnerabilities.

Fear the EAR: Discovering and Mitigating Execution After Redirect Vulnerabilities

Adam Doupe, Bryce Boe, Christopher Kruegel, and Giovanni Vigna
University of California, Santa Barbara
{adoue, bboe, chris, vigna}@cs.ucsb.edu

ABSTRACT

The complexity of modern web applications makes it difficult for developers to fully understand the security implications of their code. Attackers exploit the resulting security vulnerabilities to gain unauthorized access to the web application environment. Previous research into web application vulnerabilities has mostly focused on input validation flaws, such as cross-site scripting and SQL injection, while logic flaws have received comparably less attention.

In this paper, we present a comprehensive study of a relatively unknown logic flaw in web applications, which we call Execution After Redirect, or EAR. A web application developer can introduce an EAR by calling a redirect method under the assumption that execution will halt. A vulnerability occurs when server-side execution continues after the developer's intended halting point, which can lead to broken/insufficient access controls and information leakage. We start with an analysis of how susceptible applications written in nine web frameworks are to EAR vulnerabilities. We then discuss the results from the EAR challenge contained within the 2010 International Capture the Flag Competition. Finally, we present an open-source, white-box, static analysis tool to detect EARs in Ruby on Rails web applications. This tool found 3,944 EAR instances in 18,127 open-source applications. Finally, we describe an approach to prevent EARs in web frameworks.

Categories and Subject Descriptors

D.2.5 [Testing and Debugging]

General Terms

Security

Keywords

static analysis, web applications, execution after redirect

1. INTRODUCTION

An increasing number of services are being run online. For example, banking, shopping, social media, and entertainment are all available online. The increasing amount of sensitive data in applications has attracted the attention of cybercriminals who break into systems to steal valuable info such as passwords, credit card numbers, social security numbers, and bank account credentials.

Attackers use a variety of vulnerabilities to compromise applications. In 2008, Albert Gonzalez was later convicted of stealing 40 million credit cards from major corporate retailers, by writing SQL injection attacks [20, 30]. Another common vulnerability is cross-site scripting (XSS), is the second highest-ranked OWASP top ten security risks for web application injection attacks like SQL injection [29]. The OWASP top ten security risks for web application injection attacks like SQL injection [29]. The OWASP top ten security risks for web application injection attacks like SQL injection [29]. The OWASP top ten security risks for web application injection attacks like SQL injection [29].

In this paper, we present an in-depth study of a real-world web application logic flaw: one we call Execution After Redirect (EAR). An EAR occurs when a developer's misunderstanding of how the web framework operates. In the normal workflow of a web application, a user sends a request to the web application which receives this request, performs server-side processing, and returns an HTTP response. The HTTP response can be a notification that the user browser should look elsewhere for the response. In this case, the web application sets response codes to 301, 302, 303, or 307, and a Location header [32]. These response codes instruct the browser to look for the resource originally requested at the specified by the web application in the HTTP header [31]. This process is known as redirect: the application redirects the user to another resource. Intuitively, one assumes that a redirect should

Toward Black-Box Detection of Logic Flaws in Web Applications

Giancarlo Pellegrino
EURECOM, France
SAP Product Security Research, France
giancarlo.pellegrino@eurecom.fr

Davide Balzarotti
EURECOM, France
davide.balzarotti@eurecom.fr

Abstract—Web applications play a very important role in many critical areas, including online banking, health care, and personal communication. This, combined with the limited security training of many web developers, makes web applications one of the most common targets for attackers.

In the past, researchers have proposed a large number of white- and black-box techniques to test web applications for the presence of several classes of vulnerabilities. However, traditional approaches focus mostly on the detection of input validation flaws, such as SQL injection and cross-site scripting. Unfortunately, logic vulnerabilities specific to particular applications remain outside the scope of most of the existing tools and still need to be discovered by manual inspection.

In this paper we propose a novel black-box technique to detect logic vulnerabilities in web applications. Our approach is based on the automatic identification of a number of behavioral patterns starting from few network traces in which users interact with a certain application. Based on the extracted model, we then generate targeted test cases following a number of common attack scenarios.

We applied our prototype to seven real world E-commerce web applications, discovering ten very severe and previously-unknown logic vulnerabilities.

1. INTRODUCTION

Web applications play a very important role in many critical areas, and are currently trusted by billions of users to perform financial transactions, store personal information, and communicate with their friends. Unfortunately, this makes web applications one of the primary targets for attackers interested in a wide range of malicious activities.

To mitigate the existing threats, researchers have proposed a large number of techniques to automatically test web applications for the presence of several classes of vulnerabilities. Existing solutions span from black-box fuzzers and pentesting

tools to static analysis systems that parse the source code of an application looking for well-defined vulnerability patterns. However, traditional approaches focus mostly on the detection of input validation flaws, such as SQL injection and cross-site scripting. To date, more subtle vulnerabilities specific to the logic of a particular application are still discovered by manual inspection [33].

Logic vulnerabilities still lack a formal definition, but, in general, they are often the consequence of an insufficient validation of the business process of a web application. The resulting violations may involve both the control plane (i.e., the navigation between different pages) and the data plane (i.e., the data flow that links together parameters of different pages). In the first case, the root cause is the fact that the application fails to properly enforce the sequence of actions performed by the user. For example, an application may not require a user to log in as administrator to change the database settings (authentication bypass), or it may not check that all the steps in the checkout process of a shopping cart are executed in the right order. Logic errors involving the data flow of the application are caused instead by failing to enforce that the user cannot tamper with certain values that propagate between different HTTP requests. As a result, an attacker can try to replay expired authentication tokens, or mix together the values obtained by running several parallel sessions of the same web application.

Formal specifications describing the evolution of the internal state and of the expected user behavior are almost never available for web applications. This lack of documentation makes it very hard to find logic vulnerabilities. For example, while being able to add several items the same product to a shopping cart is a common feature, being able to add several items the same discount code is likely a logic vulnerability. A human can easily understand the difference between these two scenarios, but for an automated scanner without the proper application model it is very hard to tell the two behaviors apart.

Why Logic Bugs?

As code moves more and more to memory-safe languages (ala TRACTOR and advanced LLMs), traditional (memory) bugs are not possible.

But logic bugs are!

Example.

In 2025, Ubuntu switched from the C-based coreutils to the rust-based utils...
... and had them audited, resulting in **44** CVEs in April 2026.

But now we're done?

Why Logic Bugs?

As code moves more and more to memory-safe languages (ala TRACTOR and advanced LLMs), traditional (memory) bugs are not possible.

But logic bugs are!

Example.

In 2025, Ubuntu switched from the C-based coreutils to the rust-based utils...
... and had them audited, resulting in **44** CVEs in April 2026.

But now we're done?



We'll write secure Rust!



We'll audit code as we port it!



There's not that much code to port anyways.



LLMs will write secure Rust!



A Rust implementation of Android's Binder

This article brought to you by LWN subscribers

Subscribers to LWN.net made this article — and everything that surrounds it — possible. If you appreciate our content, please [buy a subscription](#) and make the next set of articles possible.

By **Jonathan Corbet**

November 30, 2023

[LPC](#)

The Android system was once famous for extensive, out-of-tree kernel enhancements. Many of those have been eliminated or upstreamed over the years, bringing Android much closer to the mainline kernel. One significant component in the "upstreamed" category is Binder, an interprocess communication mechanism that is used only by Android. There are a number of factors that make Binder a good candidate for rewriting in the Rust language; at the [2023 Linux Plumbers Conference](#), Carlos Llamas and Alice Ryhl described the motivation


behind and implementation of a rewrite of Binder in Rust.

ASU ran an
experiment, giving
every employee
infinite Codex
usage...

SafeLibs BETA

SafeLibs builds memory-safe (Rust) reimplementations of critical load-bearing C/C++ libraries used throughout open source infrastructure, while attempting to preserve drop-in compatibility at compile-time and runtime.

SafeLibs: Translating All C (load bearing libraries) To Rust



SafeLibs: Translating All C (load bearing libraries) To Rust



Key Insight

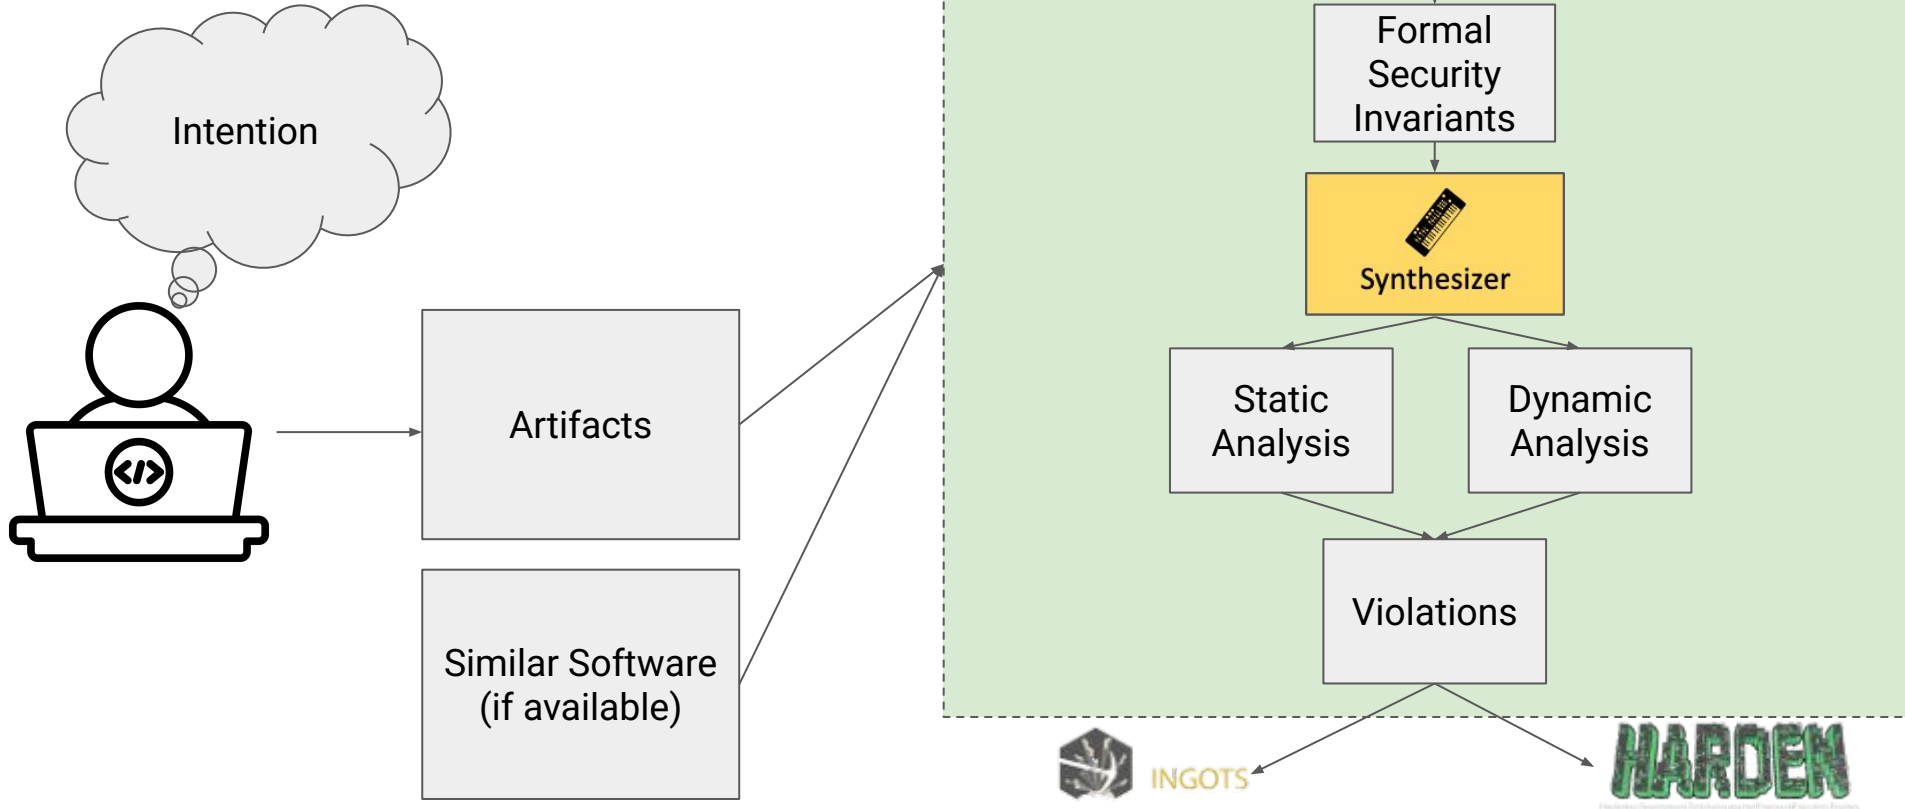
Application context is key.

Similar software shares similar intent.

LLMs contain *world knowledge* that can help identify *expected* behaviors.

Synthesize invariants that must hold in the system.

Overview Architecture



Test Target

Bluetooth Authorization.

Intent Communication.

Inter-process Communication.

Permission Control.

Unexpected Directory Traversal.

... and more



Ground Truth—Android and Open Harmony

SoK: History Doesn't Repeat Itself, but Android Design-Level Vulnerabilities Rhyme in OpenHarmony

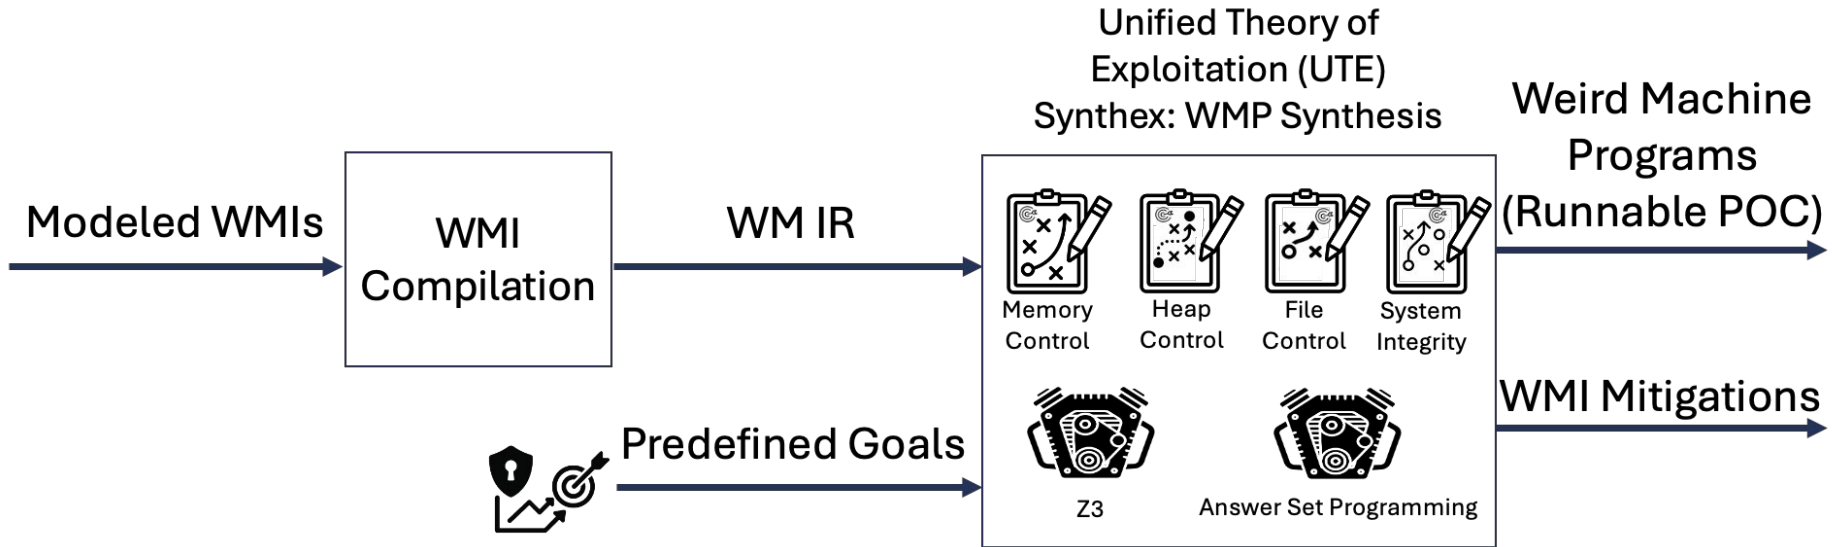
Anonymous Submission

Abstract

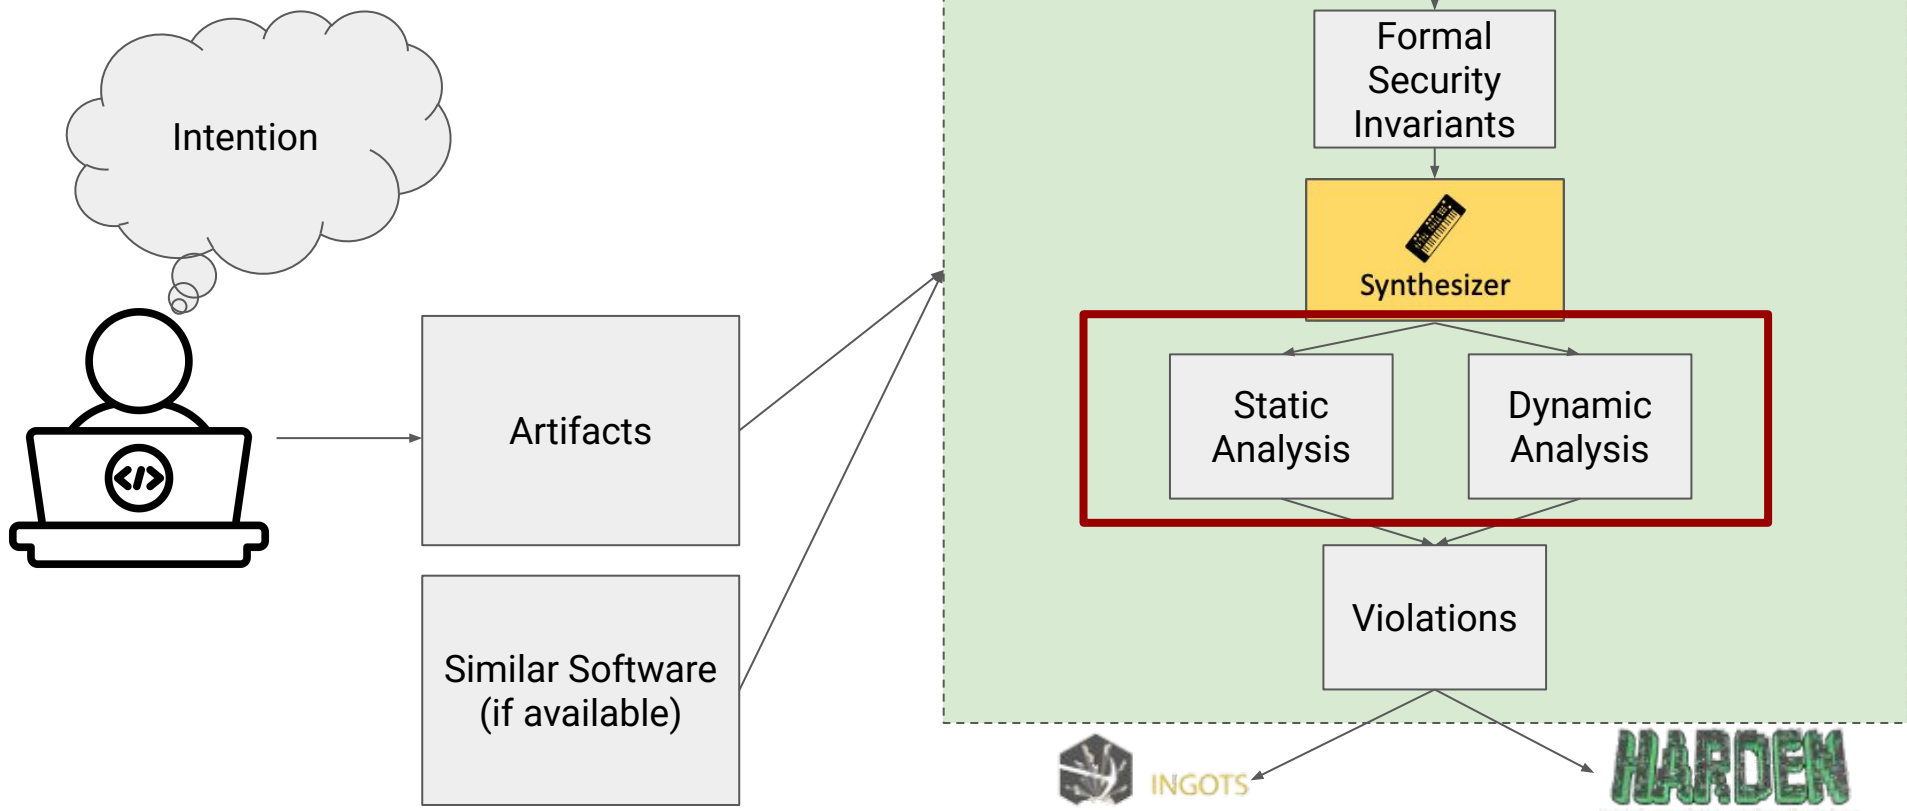
The dominant market share of Android across the world has led to significant scrutiny by security researchers. Over the years, many security issues were identified and remedied not only in its implementation but also in its design. Interestingly, academia has not further systematized these design flaws. This is an important gap, especially as many phone vendors have recently begun developing their own mobile operating systems (OSes). While they do not necessarily reuse the code of Android, they may share much of its design as they often

Recently, this fragmentation has evolved to a new stage. In 2019, due to various geopolitical, technological, and commercial factors, some Android vendors began to supersede their customizations of Android with implementations of their own mobile OSes. Many of these OSes are initially still based on Android: for example, Huawei's HarmonyOS up to version 4 was based on Android, but the developers had gradually replaced Android components with their own. Importantly, this gradual, Ship of Theseus-style replacement implies a similar design to maintain compatibility with the yet-untranslated Android components. As one illustration, an *Intent* is an object

UTE—Universal Theory of Exploitation



Status

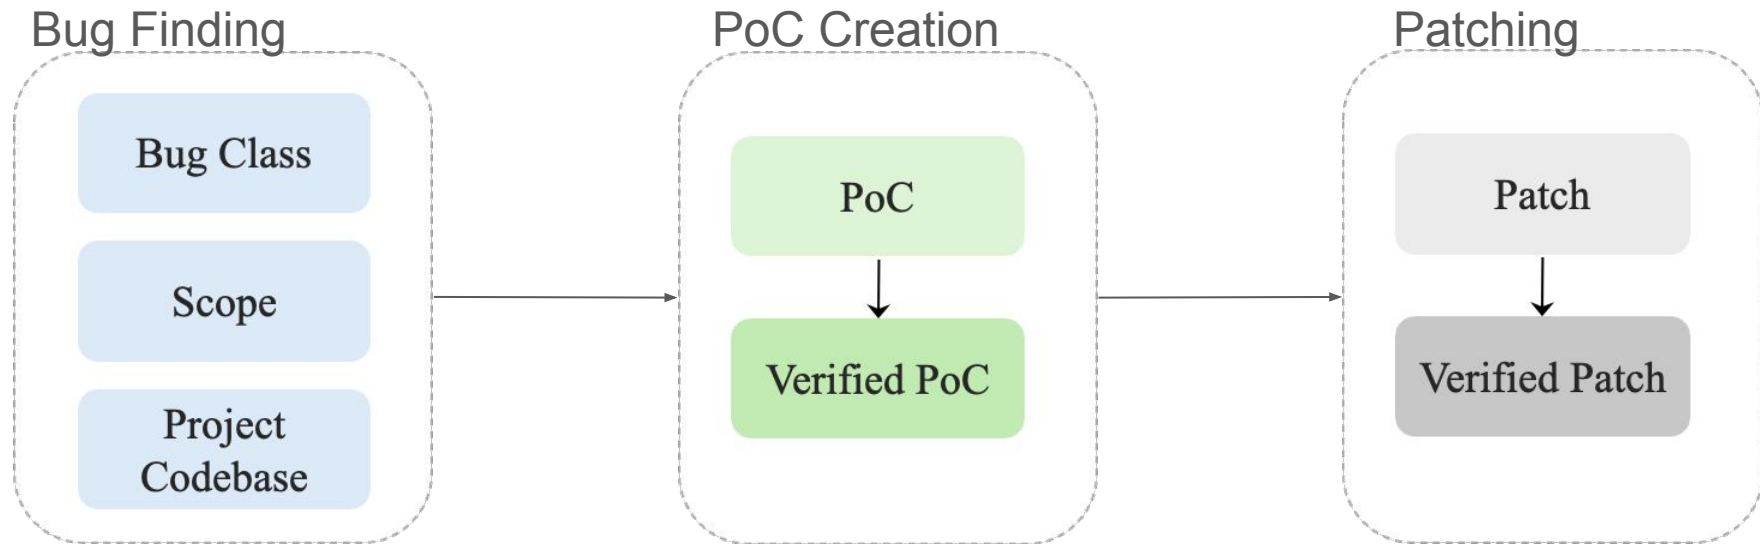





That boy is our last hope.

No. There is another.

Pure LLM Vulnerability Discovery



Can LLMs find Logic Bugs?



CVE-2026-31431 100% reliable every distro since 2017 container escape primitive 732 bytes
found by [Xint Code](#)

Copy Fail

Most Linux LPEs need a race window or a kernel-specific offset.

Copy Fail is a **straight-line logic flaw** — it needs neither.

The same **732-byte** Python script roots every Linux distribution shipped since 2017.

One logic bug in `authencsn`, chained through `AF_ALG` and `splICE()` into a

Target



Logic Bugs in XNU

Analysis found many bugs, still reporting / triaging.

Severity varies (many latent bugs).

No copy.fail (yet).

Double chroot double the fun?

XNU xnu-12377.101.15 change added support for **chroot** inside **chroot**:

```
/*
```

```
* If the process is already running with an altered root directory, then  
* the process's current directory is changed to the same new root  
* directory.  
*/
```

However, threads have a cwd in addition to the process, and that is used when opening files...

Hands Off My MAC

The MAC (Mandatory Access Control) framework in XNU is a kernel-level policy plug-in system that allows kernel extensions to control access.

Two functions use `MAC_PERFORM(...)` macro instead of `MAC_CHECK(...)`.

`MAC_PERFORM` does not set `error`...

```
if ((error = mac_vnode_check_dataprotect_set(ctx, vp,  
                                             &dataprotect_class))) {  
    vnode_put(vp);  
    goto outdrop;  
}
```

An Immovable Right Meets a Logic Bug

Mach part of XNU using Mach Messages to do IPC. (Much more complex than fds.)

`IPC_MOVE_POLICY_NEVER` is documented as a hard invariant where a receive right cannot be moved.

But this can be bypassed by arming port-destroyed while movable, then mark immovable via the guard path, and the kernel ships the receive right to the notification receiver!

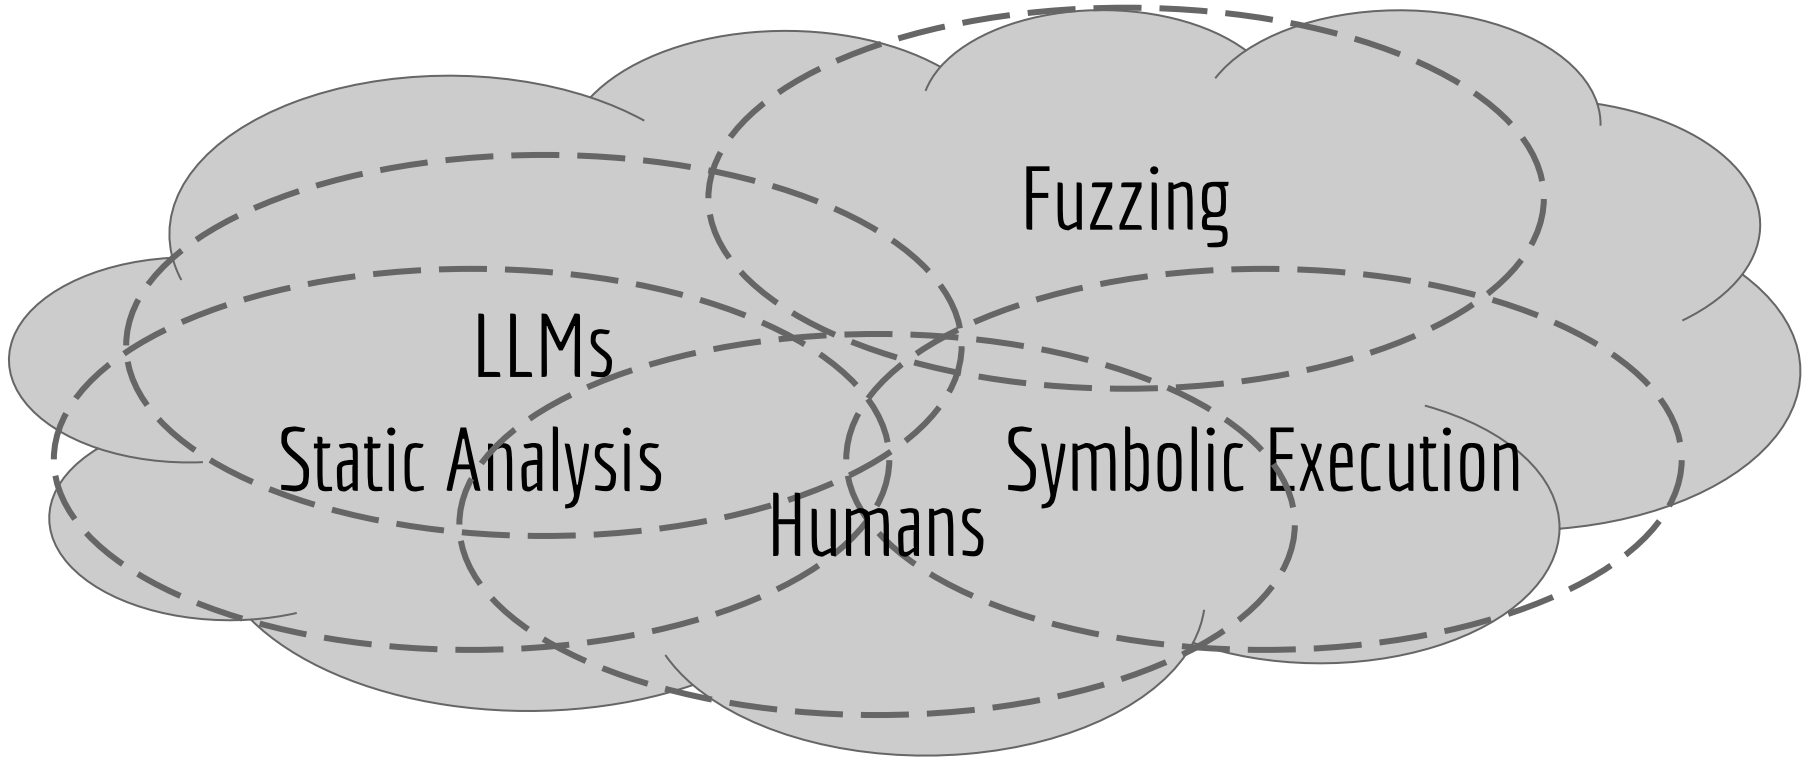
Logic Bugs as Weird Machine Instructions

One their own, none of these bugs directly seem to have a security impact.

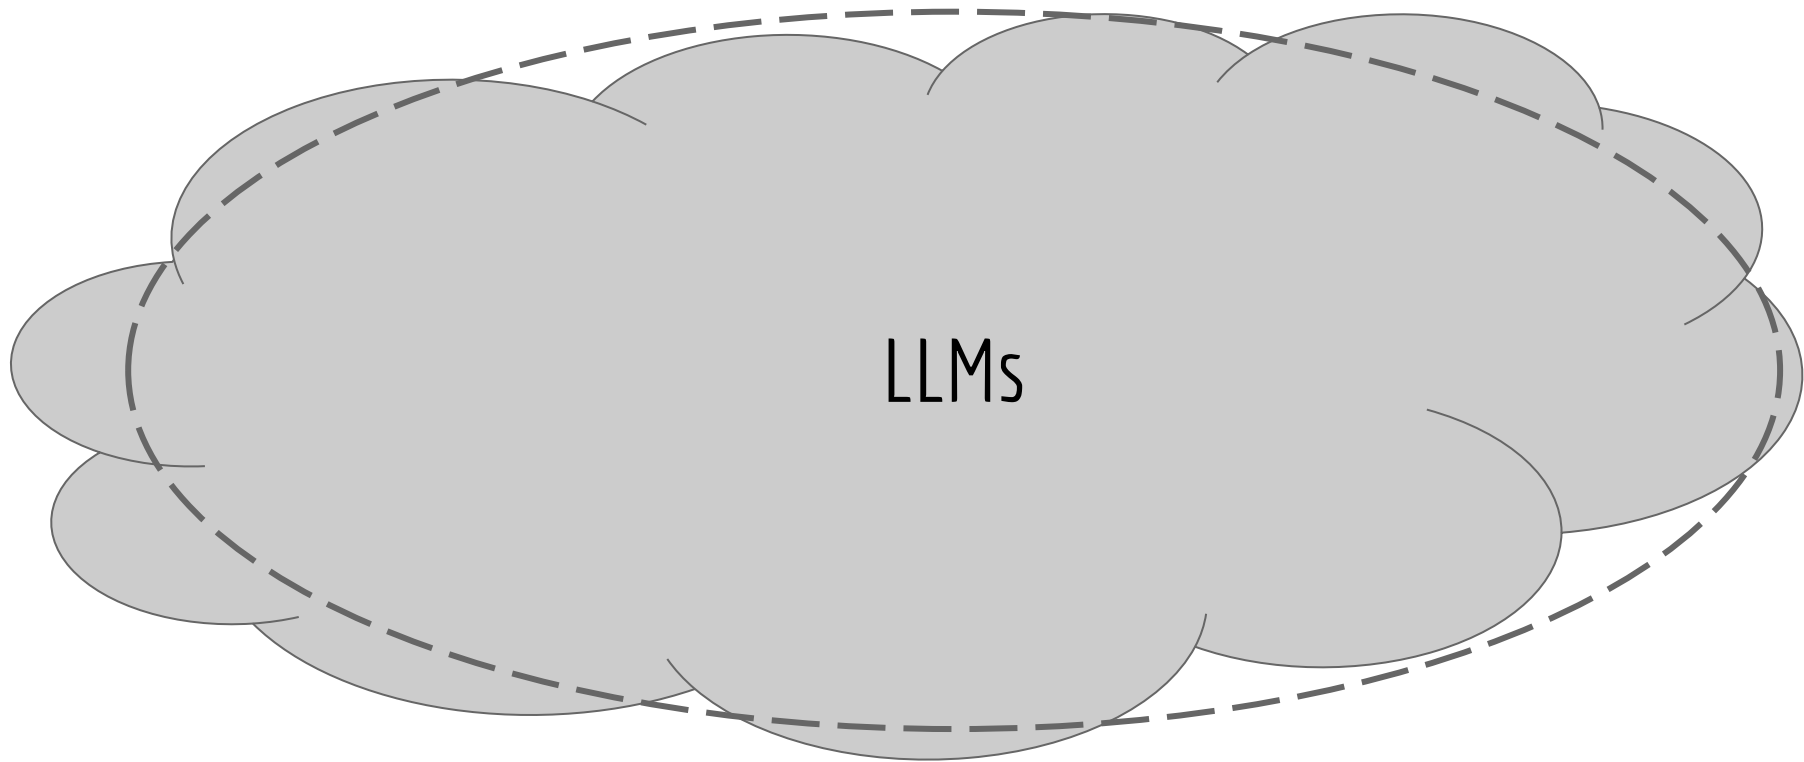
The system enters a weird program state (logical invariants are violated), could be chained and lead to future vulnerabilities.

This was the case for copy.fail, in which three bugs with seemingly no security impact came together for an epic vulnerability...


Why Two Approaches?



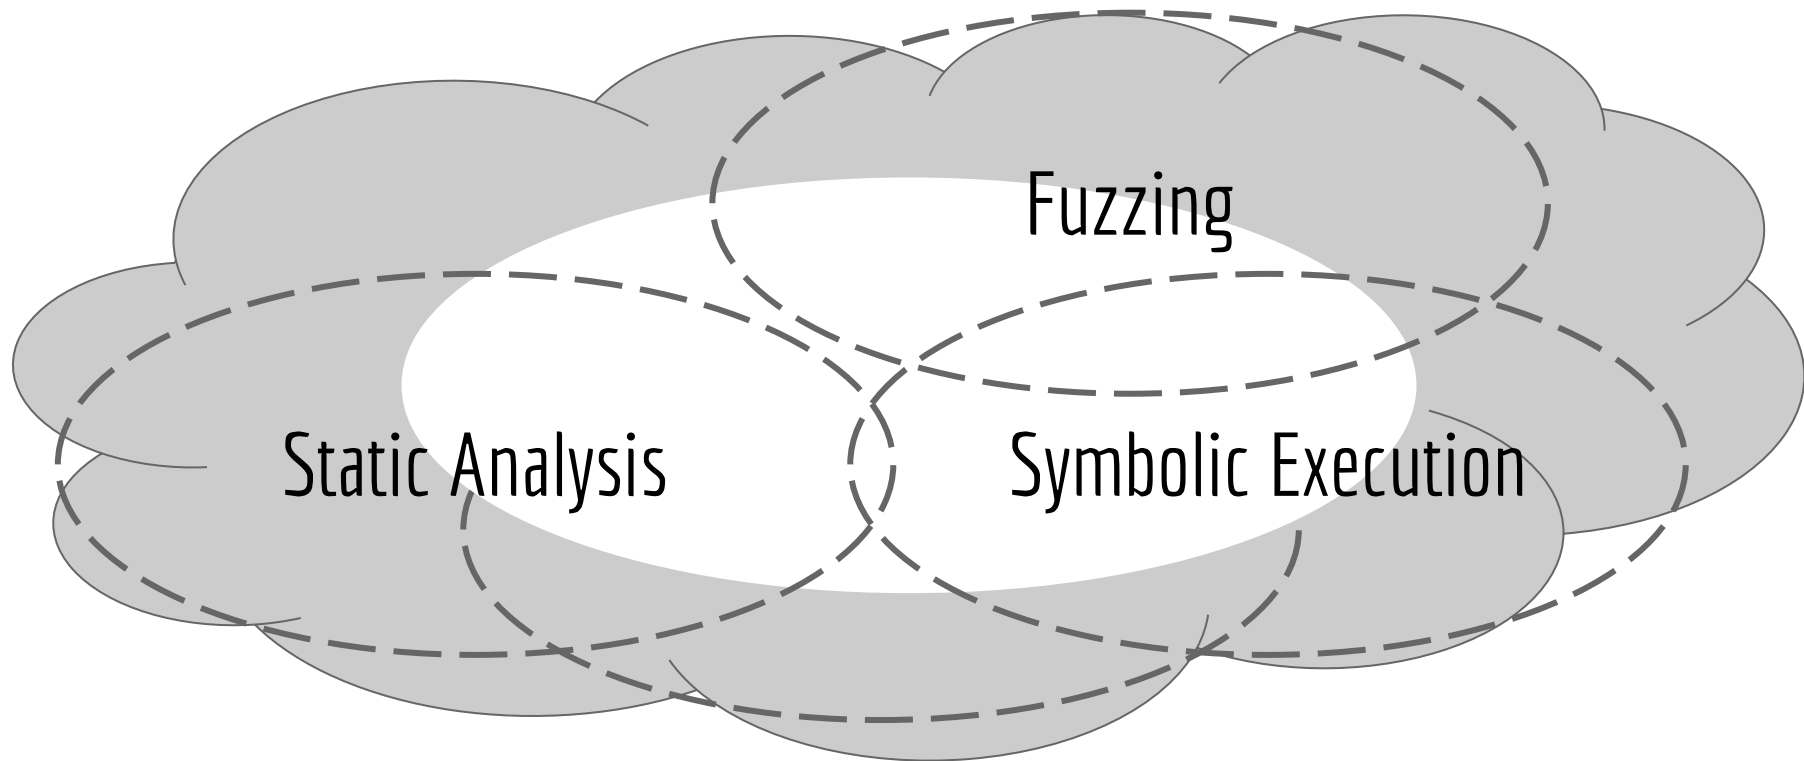
LLMs?



LLMs?



Short-term vs. Long-term



Remember SafeLibs?

Rust reimplementation re-introduced several logic bugs!

```
libgcrypt – CVE-2024-2236 – RSA PKCS#1 v1.5 Bleichenbacher oracle
libgcrypt – CVE-2018-6829 – ElGamal m=0 leaks b=0
libxml – CVE-2024-40896 – XInclude file:// XXE under default parser flags
libtiff – CVE-2023-6277 – StripByteCounts vec![0u8; N] aborts on OOM
libtiff – CVE-2017-12944 – XMLPacket count vec![0u8; N] aborts on OOM
libtiff – CVE-2018-5784 – StripOffsets count vec![0u8; N] aborts on OOM
libvips – CVE-2026-3146 – Matrix loader Vec::with_capacity(2^60) aborts
libpng – CVE-2014-0333 – Progressive reader never fires end_cb on zero-length IDAT
libjson – CVE-2013-6371 – PERLLIKE hash collisions → 260× DoS (opt-in)
libc6 – CVE-2010-4051 – regcomp bounded-repetition unvalidated → OOM
libc6 – CVE-2010-4756 – glob unmatched-pattern exponential walk
```

Remember SafeLibs?

Rust reimplementation re-introduced several logic bugs!

```
libcrypt - C  
libcrypt - C  
libxml - C  
libtiff - C  
libtiff - C  
libtiff - C  
libvips - C  
libpng - C  
libjson - C  
libc6 - C  
libc6 - C
```



Live Long and Prosper... Free of Logic Bugs

