


Fuzzing Rust Smart Contracts

Writing a bug printer engine from scratch

Kevin Valerio <kevin@srlabs.de>, Daniel Schmidt <schmidt@srlabs.de>



Security
Research
Labs



A **\$125 million** mistake —
caused by a
simple,
overlooked flaw

Nice to meet you



Kevin Valerio

Security Engineer
at SRLabs 

Background in
pentesting and
Web3 security



Daniel Schmidt

Security Researcher
at SRLabs 

Background in
protocol and virtual
machine security



Agenda

1. Overview

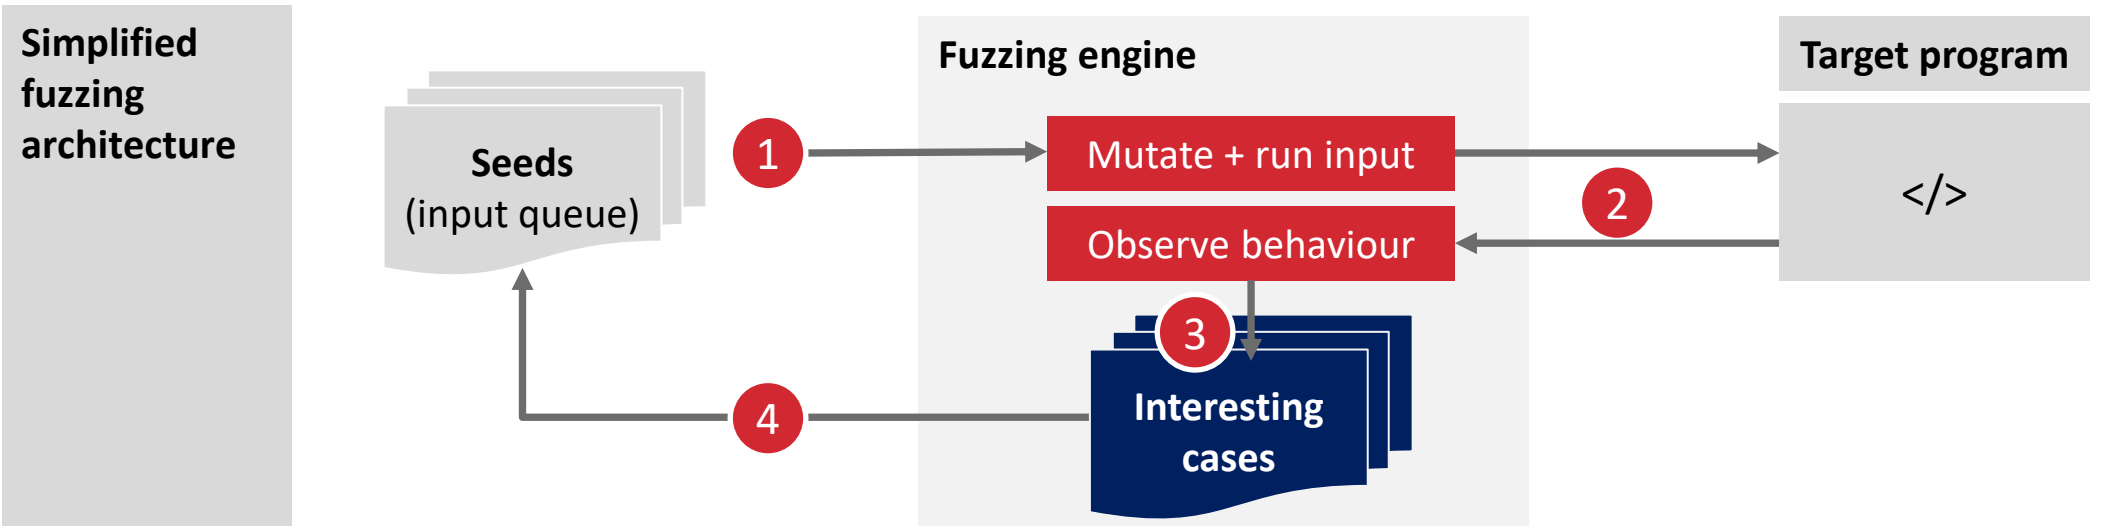
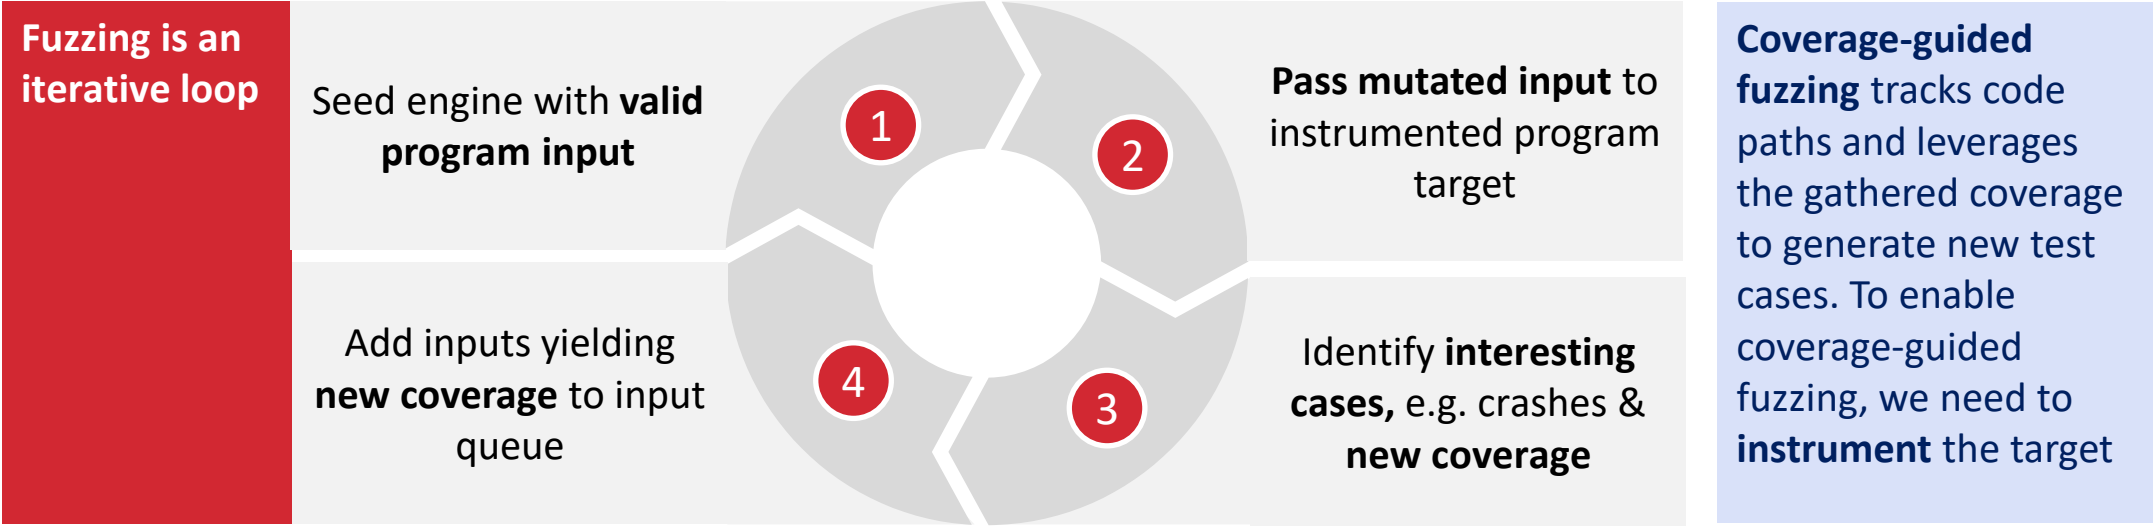
▶ 2. Background

3. Challenges

4. Solutions

5. Success

Fuzzing identifies vulnerabilities via mutating valid program inputs



Instrument the target by injecting callbacks to enable coverage-guided fuzzing

Target instrumentation and coverage callbacks

- 1 Identification of basic blocks
- 2 Insertion of fuzzer callback at every basic block
- 3 Callbacks write to coverage map during execution
- 4 Evaluation of coverage by the fuzzing engine



Instrumentation example

Target Code

```
void parse_input(char *input) {  
    if (input[0] == 'F') {  
        if (input[1] == 'U') {  
            if (input[2] == 'Z') {  
                if (input[3] == 'Z') {  
                    // Crash here  
                }  
            }  
        }  
    }  
}
```



Instrumented Code

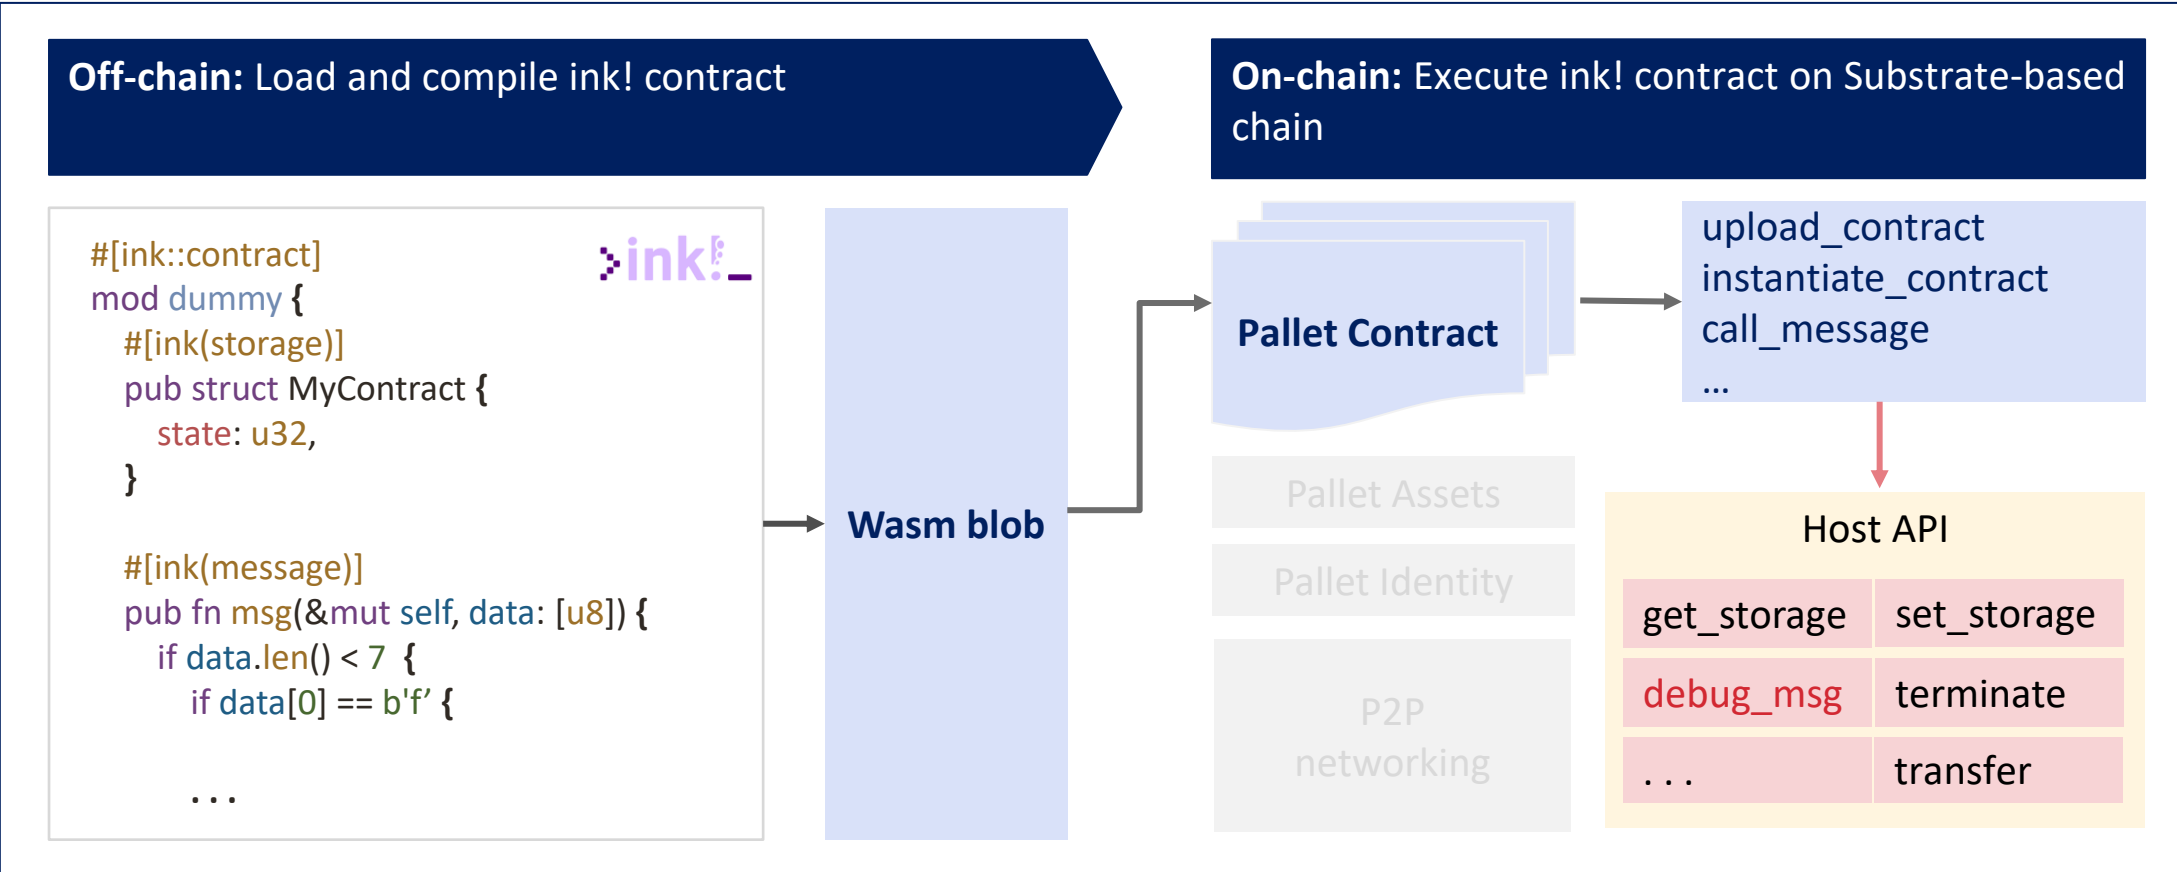
```
void parse_input(char *input) {  
    if (input[0] == 'F') {  
        __sanitizer_cov_trace_pc()  
        if (input[1] == 'U') {  
            __sanitizer_cov_trace_pc()  
            if (input[2] == 'Z') {  
                __sanitizer_cov_trace_pc()  
                if (input[3] == 'Z') {  
                    __sanitizer_cov_trace_pc()  
                    // Crash here  
                }  
            }  
        }  
    }  
}
```

ink! smart-contracts are permissionless programmable extensions deployed on the blockchain

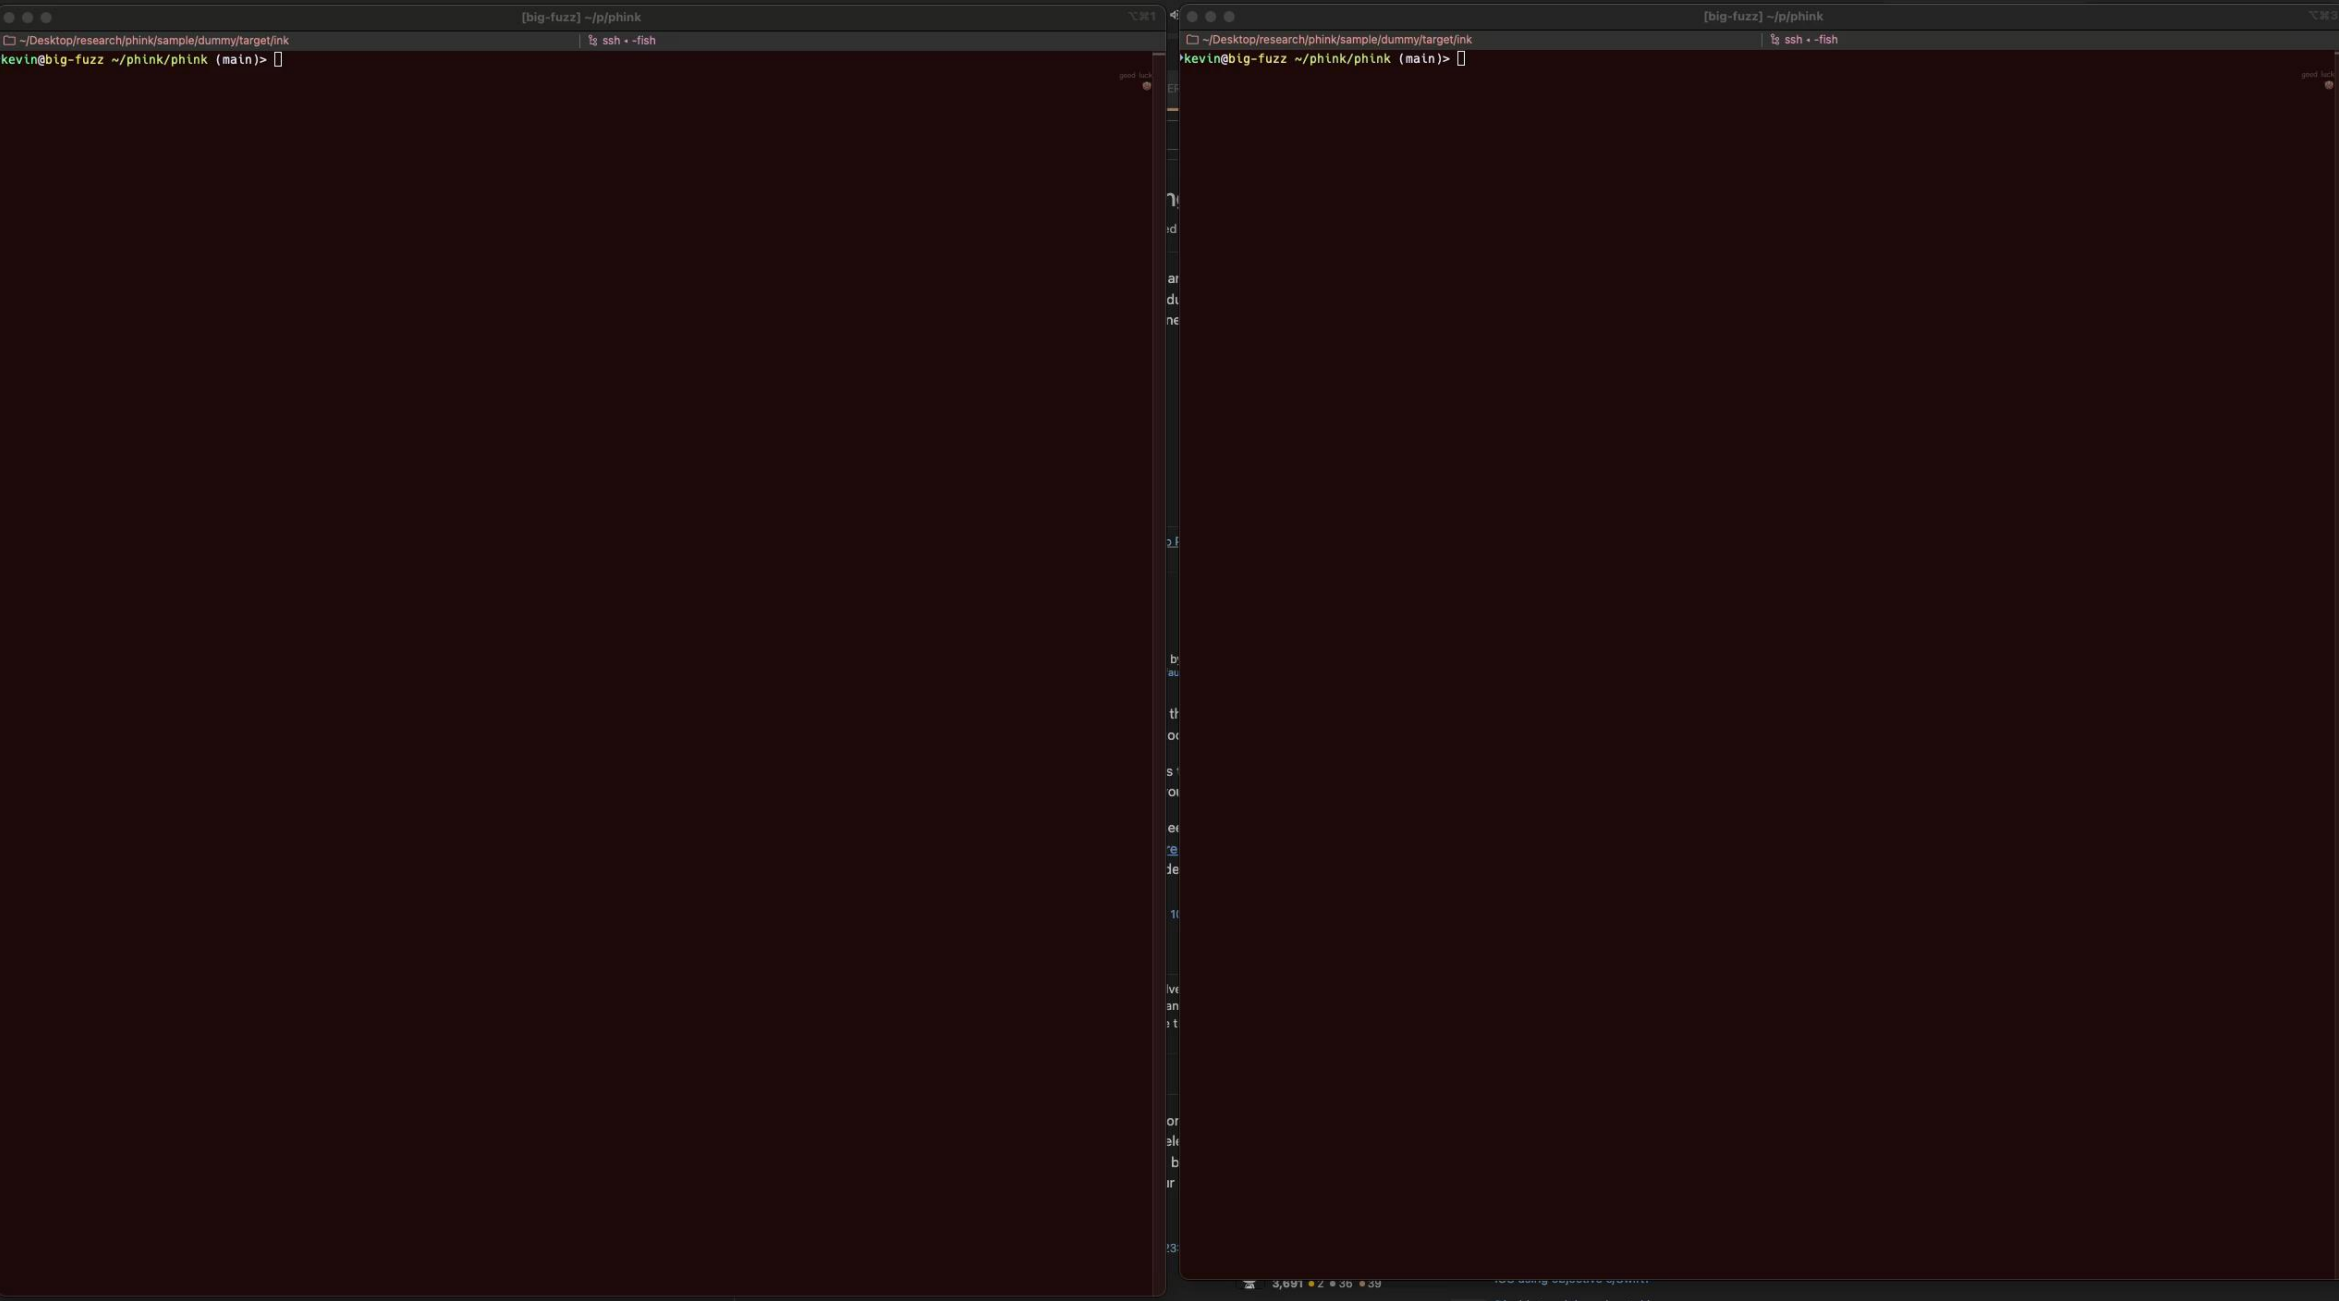
Description

- Smart contracts is **permissionless code** running inside the blockchain
- **ink! is a programming language** for smart contracts within the Polkadot ecosystem
- Being able to execute **cross-chain transactions** from ink! makes it special within the ecosystem of smart contracts

Architecture



We present Phink, a coverage guided fuzzer for ink! smart contracts



Agenda

1. Overview

2. Background

▶ 3. Challenges

4. Solutions

5. Success

Several challenges have been identified during the creation of phink

Challenge	Details
1 Execution and instrumentation barriers	<ul style="list-style-type: none">▪ ink! contracts run in a VM, preventing direct instrumentation▪ Standard fuzzers struggle to track execution paths in sandboxed environment
2 Initial seed generation	<ul style="list-style-type: none">▪ Fuzzing campaigns need initial seeds so that they do not solely rely on random chance▪ Creating initial seeds automatically is desirable
3 Stateful execution and on-chain dependencies	<ul style="list-style-type: none">▪ Smart contracts interact with on-chain data and previous state▪ Ensuring meaningful multi-call transactions during fuzzing is complex
4 Coverage and feedback limitations	<ul style="list-style-type: none">▪ Generating coverage reports is difficult but crucial for optimizing fuzzing campaigns▪ Limited visibility into how much of the contract is being tested

① Coverage-guided fuzzing on VMs is challenging due to execution abstraction

I

Instrumenting ink!

Ink! compiler has constraints. Typically, code is instrumented by compilers (e.g., *afl-clang*). Ink! uses its custom compiler, which lacks native instrumentation

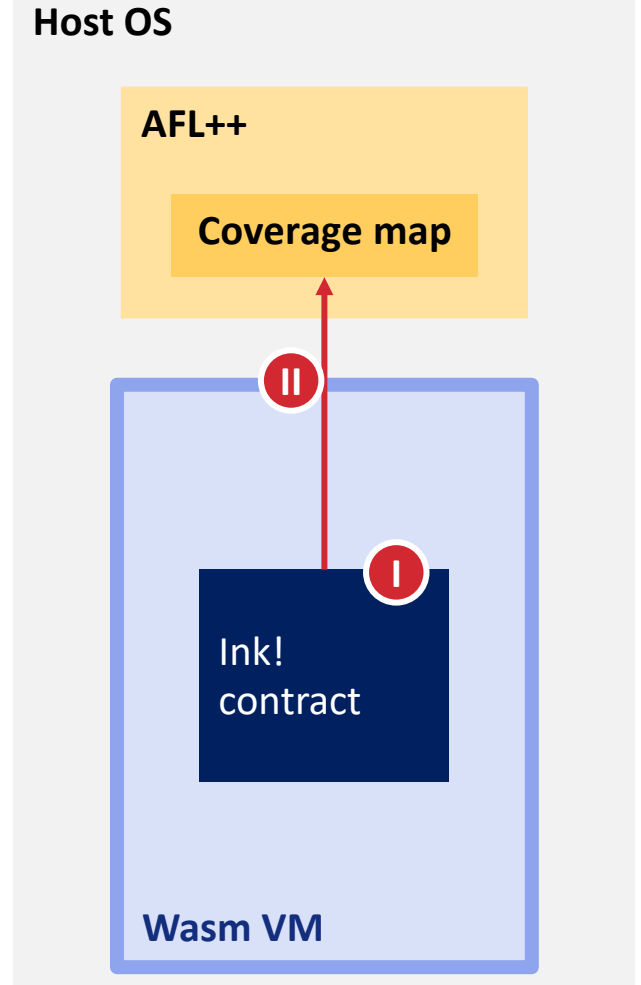
Instrumentation requires support. A version of the ink! compiler must be forked and maintained, or a PR submitted to the ink! compiler, both requiring ongoing maintenance

II

Passing through VM Sandbox

Sandbox restrictions. The Wasm Virtual Machine operates in a sandbox, making it challenging to pass information outside of the VM

Escape the sandbox to transmit coverage to AFL++. We need a way to transmit coverage beyond the sandbox and store it in AFL++'s coverage map



2 Creating effective initial seeds for a fuzzing campaign is difficult

I Impact of Initial Seeds


Academia has shown that **selecting appropriate initial** seeds can significantly impact the success of a fuzzing campaign

II Manual Seed Generation

The manual creation of seeds is a **time-consuming process**. Therefore, automating this procedure is favourable

III Seeds Harnessing Source

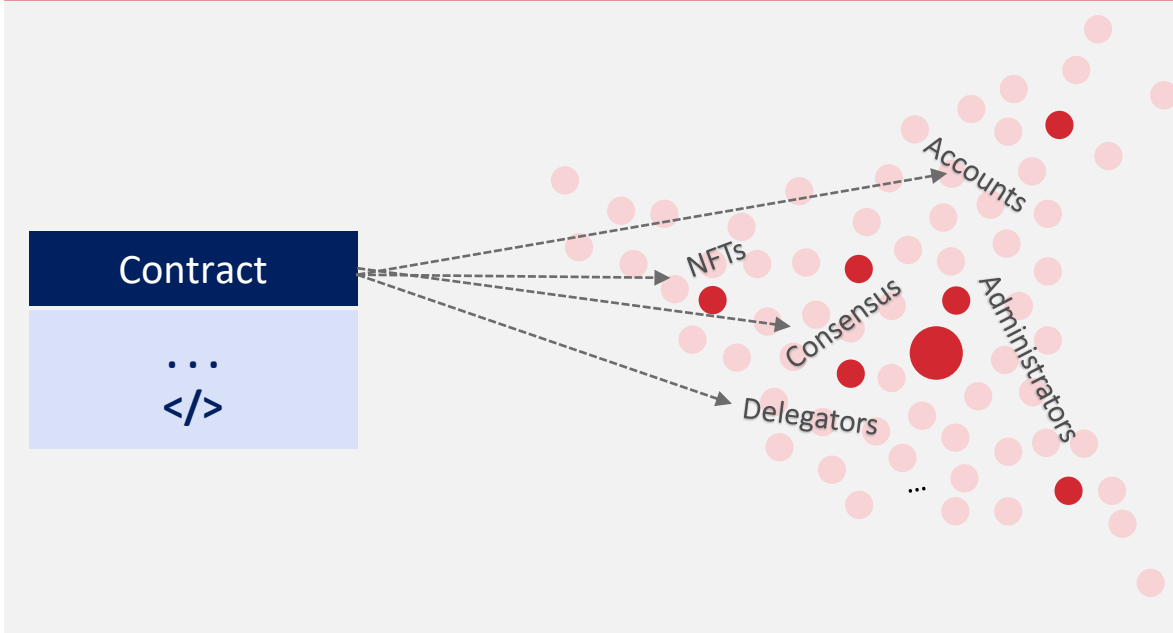
To automatically generate fuzzing seeds, we need to harness them from a **reliable source**



! An initial corpus covering more regions yields higher coverage over time

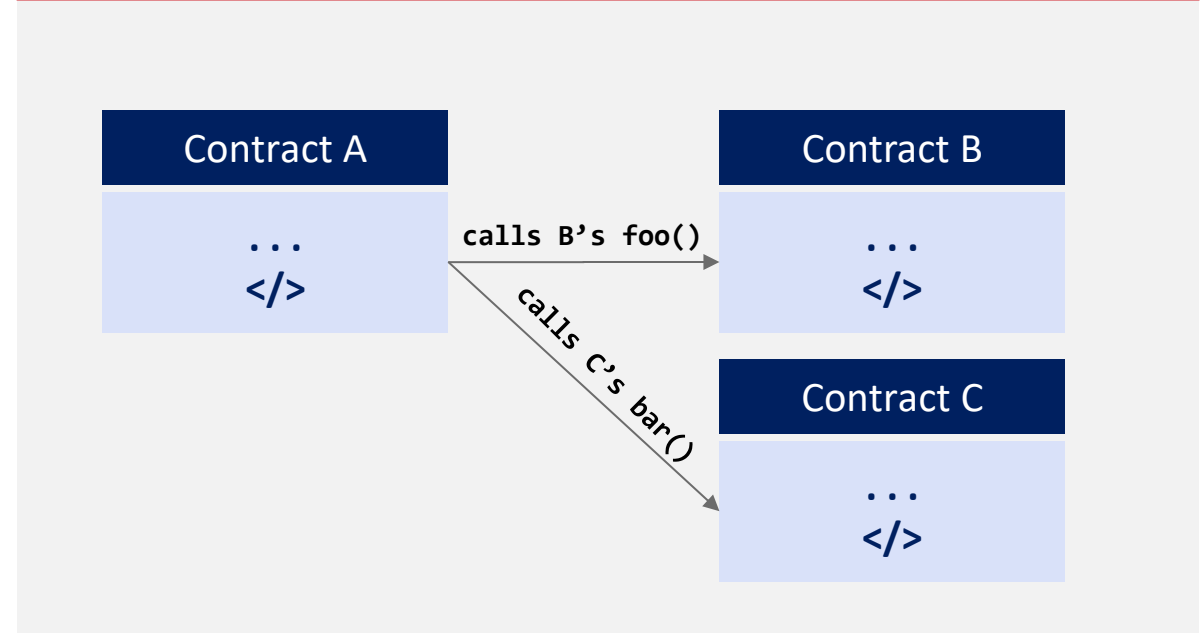
3 Bootstrapping Smart Contract with stateful execution and on-chain dependencies

I Contract requires on-chain state to work properly



? How can we supply **real-world state** data to ensure the contract functions properly?

II Contract might interact other deployed contracts



? How do we ensure that the contract can **interact** with its **contract dependencies**?

4 Generating coverage reports is crucial to have a successful fuzzing campaign

Coverage analysis is essential for discovering gaps in your fuzzing campaign, for example, through a misconfiguration

There are tools for generating coverage reports, but what if you have your own coverage system?

I The fuzzer hits this line only a **few times**; therefore, some edge cases may still be untested

II Has been hit almost as much as the function itself and can be considered **well covered**

III This line has **not been hit**, and either custom seeds or adaptation to the corpus might be required

OSS-Fuzz Coverage Report for cgif

71		/* create a new GIF */
72	4.24k	CGIF* cgif_newgif(CGIF_Config* pConfig) {
73	4.24k	FILE* pFile;
74	4.24k	CGIF* pGIF;
75	4.24k	CGIFRaw* pGIFRaw; // raw GIF stream
76	4.24k	CGIFRaw_Config rawConfig = {0};
77		// width or height cannot be zero
78	4.24k	if(!pConfig->width !pConfig->height) {
79	10	return NULL;
80	10	}
81	4.23k	pFile = NULL;
82		// open output file (if necessary)
83	4.23k	if(pConfig->path) {
84	2.13k	pFile = fopen(pConfig->path, "wb");
85	2.13k	if(pFile == NULL) {
86	0	return NULL; // error: fopen failed
87	0	}
88	2.13k	}

Agenda

-
1. Overview
 2. Background
 3. Challenges
 - ▶ 4. Solutions**
 5. Success
-

Each challenge has been addressed with solutions that will now be detailed

Challenge	Solution	Details
1 Execution and instrumentation barriers	A Custom instrumentation and coverage mapping	<ul style="list-style-type: none">▪ Phink solves instrumentation by injecting callbacks into contracts. Those callbacks serve as a direct communication from Phink to AFL++
2 Generating initial seeds	B Leverage tests for seed generation	<ul style="list-style-type: none">▪ Contract tests are leveraged to generate initial seeds▪ This provides a fully automatic and reliable method for generating initial seeds for the fuzzing campaign
3 Stateful execution and on-chain dependencies	C On-chain contract emulation and genesis state	<ul style="list-style-type: none">▪ Phink enables developers to integrate ready-to-fuzz contract dependencies and define a genesis state, creating a rich execution environment for stateful fuzzing
4 Coverage and feedback limitations	D Coverage reports	<ul style="list-style-type: none">▪ Phink's tracking of every executed statement allows to generate coverage reports which improve with monitoring and assessment of fuzzing campaigns

A Think solves instrumentation via coverage remapping and message bridging

- Think**
- I Custom instrumentation and compilation helper**
Parse smart contract code using syn lib and inject custom instrumentation, then compile
 - II Fuzz runner**
Spawn AFL++ instrumented fuzz coverage redirector and update AFL++ shared memory map utilizing a message bridge

```
Instrumented smart contract  
fn foo() {  
  ink::env::debug_println!("1")  
  let bytes = [0x4e,0x75,0x6c,0x6c,0x63,0x6f,0x6e];  
  ink::env::debug_println!("2")  
  let string = String::from_utf8_lossy(&bytes);  
  ink::env::debug_println!("3")  
  return;  
}
```

ink! WASM host including message bridge

Rust-side instrumented smart contract binary

Instrument

Compile

Run

Coverage Feedback

```
AFL++ instrumented  
fn redirect_coverage(&self, wasm_cov: &[u64]) {  
  ...  
  if wasm_cov.contains(1) {  
    redirect_edge_to_afl(1);  
  }  
  if wasm_cov.contains(2) {  
    redirect_edge_to_afl(2);  
  }  
  ...  
}
```

A Alternative approach: WASM blob instrumentation for resolving partial coverage & enabling black-box fuzzing

Phink

I Instrument the WASM blob without source-code
Parse the WASM code, search for control-flow instructions and insert callbacks

II Fuzz runner
Spawn AFL++ instrumented fuzz coverage redirector and update AFL++ shared memory map utilizing a message bridge

```
Compiled WASM contract  
i32.const 117  
i32.ne  
if ;; label = @3  
  local.get 6  
i32.ne  
  br_if 1 (;@2;)  
  i32.const 66200  
  call 16  
end
```

```
Instrumented and compiled WASM contract  
i32.const 117  
i32.ne  
i32.const 1000  
i32.const 4  
call 8  
drop  
if ;; label = @3  
  local.get 6  
i32.ne  
  br_if 1 (;@2;)  
  i32.const 66200  
  call 16  
end
```

ink! WASM host including message bridge

Deployed

WASM-side instrumented smart contract binary

Parse

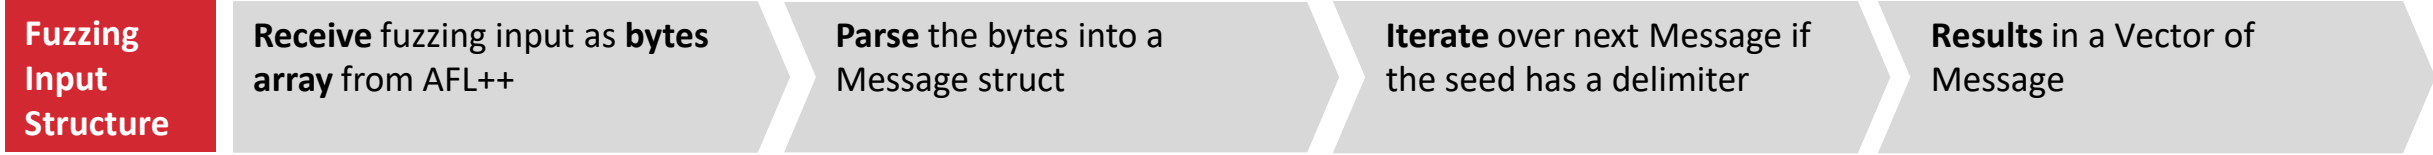
Inject callback

Run

```
AFL++ instrumented  
fn redirect_coverage(&self, wasm_cov: &[u64]) {  
  ...  
  if wasm_cov.contains(1) {  
    redirect_edge_to_afl(1);  
  }  
  if wasm_cov.contains(2) {  
    redirect_edge_to_afl(2);  
  }  
  ...  
}
```

Coverage Feedback

B Think defines a structured format for inputs, designed specifically for message execution



Message structure

```
#[derive(Debug, Clone, Serialize)]  
pub struct Message {  
    origin: Origin,  
    value_token: BalanceOf<Runtime>,  
    payload: Vec<u8>,  
}
```

u8
u128
selector
params



Multi-message structure example

Example Message 1

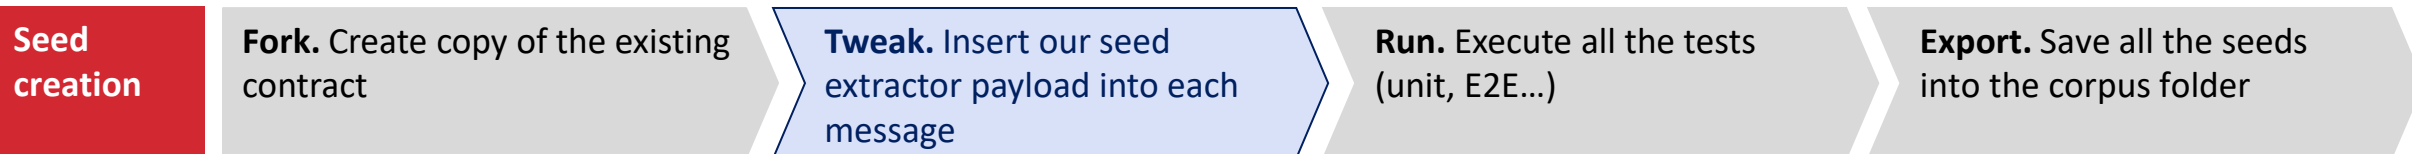
Bob	231	send_to	{ account: 5D35...x}	****
-----	-----	---------	----------------------	------

...

Example Message N

Alice	420	send_to	{ account: 5D35...x}	****
-------	-----	---------	----------------------	------

Contract tests are leveraged to generate initial seeds



Tweak phase

- 1 For each `#[ink(message)]`
- 2 **Grab** arguments + function name
- 3 Prepare a snippet that **SCALE-encode** the message selector + parameters
- 4 **Insert** that snippet at the beginning of the message

Message is tweaked to output the encoded seed when called

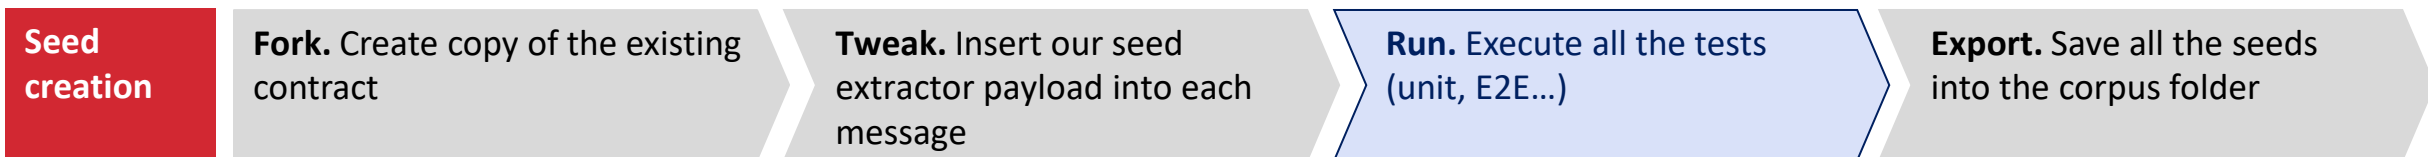
```

#[ink(message)]
pub fn crash_with_invariant(&mut self, data: String) -> Result<()> {
{
    let sel = ExecutionInput::new(selector_bytes!("crash_with_invariant"))
        .push_arg(&data);

    let encoded = scale::Encode::encode(&sel);
    ink::env::debug_println!("ENCODED_SEED={}", encoded.iter()
        .map(|byte| format!("{:02x}", byte))
        .collect::<String>());
}
// Actual message logic below
if data.len() < 7 && data.len() > 3 {
    ...
}
}
Ok(())
}
  
```

Inserted snippet

Contract tests are leveraged to generate initial seeds



Tests are executed and saved as valid seeds

```

$ cargo run -- generate-seed sample/dummy/

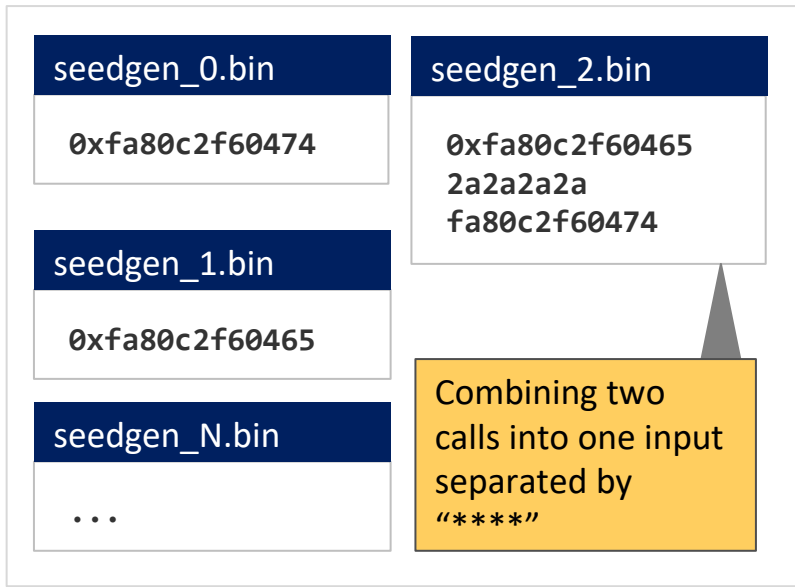
running 3 tests
test dummy::e2e_tests::it_works ... ok
test dummy::tests::for_seedgen ... ok
test dummy::tests::new_works ... ok

ENCODED_SEED=fa80c2f60474
ENCODED_SEED=fa80c2f60465

Writing bytes 0xfa80c2f60474 to
`output/phink/corpus/seedgen_0.bin`
Writing bytes 0xfa80c2f60465 to
`output/phink/corpus/seedgen_1.bin`
  
```

Seed
List of messages

Corpus directory



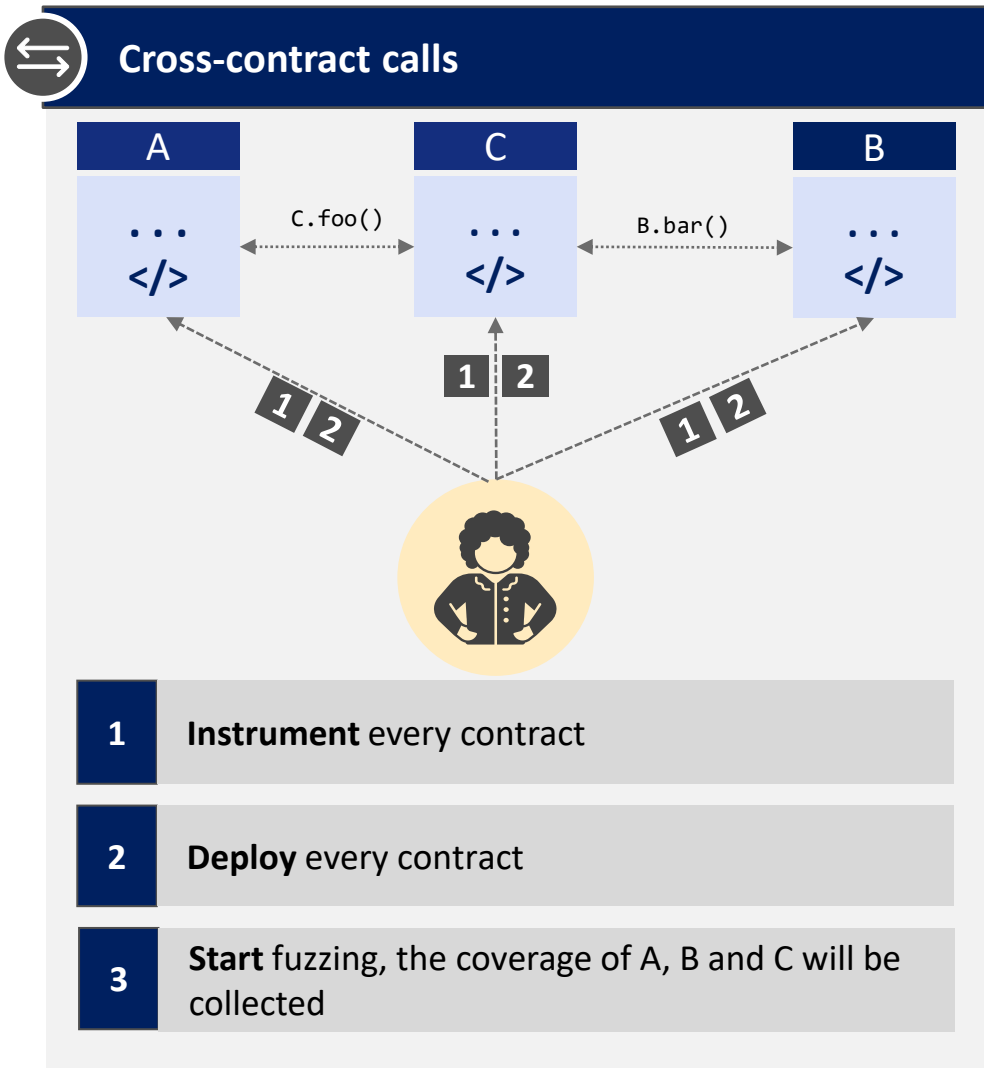
B

Contract tests are leveraged to generate initial seeds

```
[big-fuzz] ~/p/phink
~/Desktop/research/phink/sample/dummy/target/fink
kevin@big-fuzz ~/phink/phink (main) [SIGINT] > 
```

good luck





Mocking on-chain state

Developers can insert a mocked environment

```

impl DeveloperPreferences for Preferences {
  fn runtime_storage() -> Storage {
    let storage = RuntimeGenesisConfig {
      balances: BalancesConfig {
        // Lot of money for Alice, Bob..
        balances: (0..u8::MAX)
          .map(|i| [i; 32].into())
          .map(|k| (k, 50000))
          .collect(),
      },
    }.build_storage()
  }
  ...
}

```

0x00..001 (Alice) Balance: 50000

0x00..002 (Bob) Balance: 50000

0x00..003 (Charlie) Balance: 50000

D Users can create coverage reports for their contract

Think tracks every executed statement and allows to generate coverage reports

I Generating a .cov file

- 1 For each seed in corpus
- 2 Run them with the harness
- 3 Append the reached coverage into traces.cov

II Parsing and generating HTML

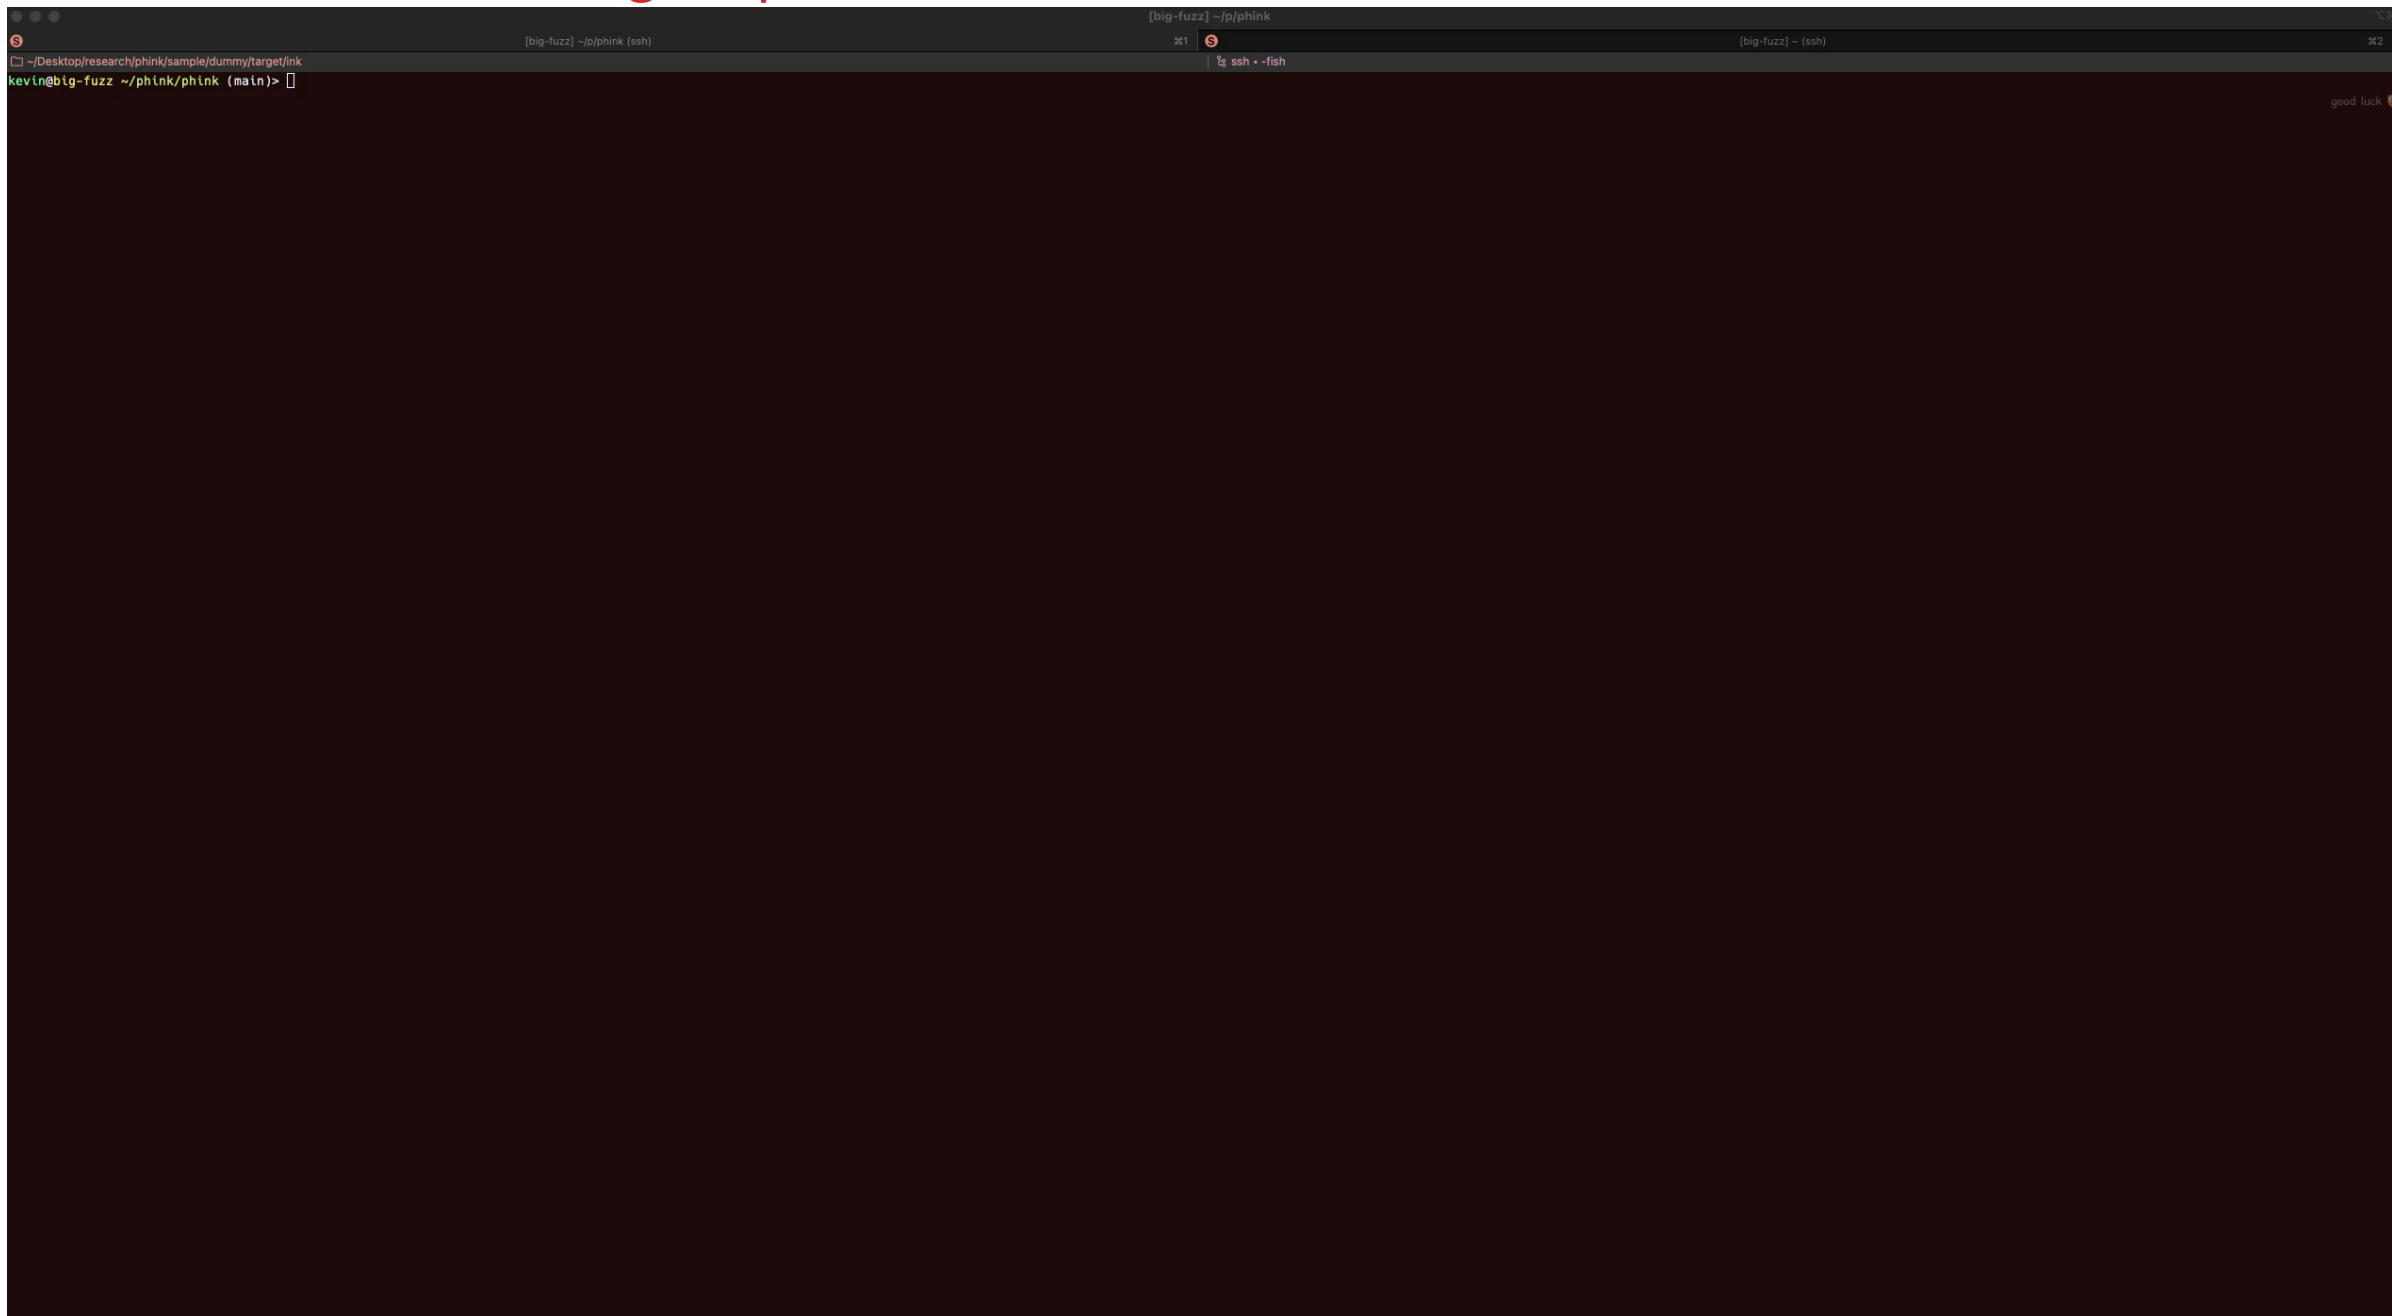
- 1 Copy the Rust files of the contract into HTML
- 2 Parse traces.cov
- 3 For each trace in traces.cov
- 4 Change the executed line's color to green

Coverage report sample

```
1 | use crate::modules::govern::traits::UnstakePeriodChanged;
2 | use ink::{
3 |     env::DefaultEnvironment,
4 |     primitives::AccountId,
5 | };
6 | pub use pendzl::contracts::general_vest::GeneralVestRef;
7 | use pendzl::traits::Timestamp;
8 | #[derive(Debug, Default)]
9 | #[pendzl::storage_item]
10 | pub struct UnstakeData {
11 |     #[lazy]
12 |     general_vester: GeneralVestRef,
13 |     #[lazy]
14 |     unstake_period: Timestamp,
15 | }
16 | impl UnstakeData {
17 |     pub fn new(general_vester_address: AccountId, unstake_period: Timestamp) -> Self {
19 |         let mut instance = Self::default();
21 |         instance.set_general_vester(&general_vester_address);
23 |         instance.set_unstake_period(unstake_period);
25 |         ink::env::emit_event::<DefaultEnvironment, UnstakePeriodChanged>(UnstakePeriodChanged {
26 |             unstake_period,
27 |         });
29 |         instance
30 |     }
31 | }
32 | impl UnstakeData {
33 |     pub fn general_vester(&self) -> GeneralVestRef {
35 |         self.general_vester.get().unwrap()
36 |     }
37 |     pub fn unstake_period(&self) -> Timestamp {
39 |         self.unstake_period.get().unwrap_or_default()
40 |     }
41 |     pub fn set_general_vester(&mut self, vester: &AccountId) {
43 |         let vester: GeneralVestRef = (*vester).into();
45 |         self.general_vester.set(&vester);
46 |     }
47 |     pub fn set_unstake_period(&mut self, period: Timestamp) {
49 |         self.unstake_period.set(&period);
50 |     }
51 | }
```


D

Users can create coverage reports for their contract



```
[big-fuzz] ~/p/phink
[big-fuzz] ~/p/phink (ssh) #1
[big-fuzz] ~ (ssh) #2
good luck
kevin@big-fuzz ~/phink/phink (main) >
```




Agenda

-
1. Overview
 2. Background
 3. Challenges
 4. Solutions

 **5. Success**

Phink is now the industry standard fuzzer for ink! smart contracts



Forkless fuzzer

Phink doesn't need a fork of ink, pallet_contract, substrate, polkadot or cargo-contract

VM-agnostic

ink! contract compiled into WASM or RISC-V (newly supported) can be fuzzed by Phink

Fully coverage guided

With in-WASM instrumentation, the contract is fully instrumented on every control-flow

Blackbox fuzzing

Since Phink can instrument compiled WASM blobs, source-code is not required

