

# OSS-CRS: Bug-Finding and Remediation for the LLM era

Andrew Chin

SSLab @ Georgia Institute of Technology

Team Atlanta




OpenSSF Community Day  
NORTH AMERICA 2026

# Bug-finding

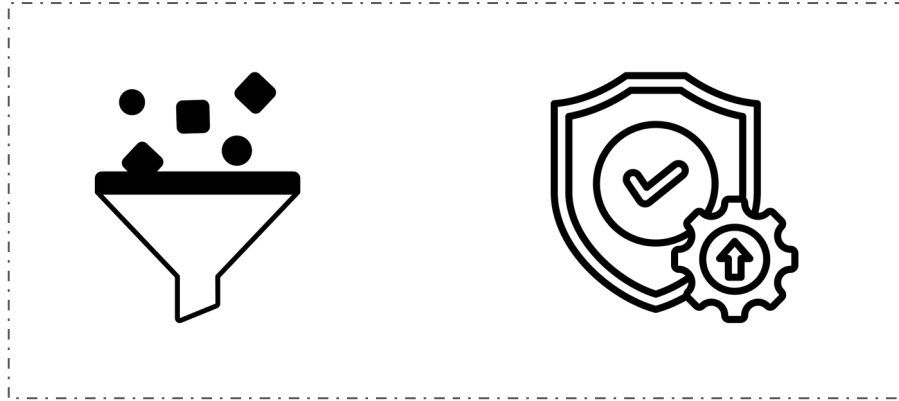



# Triaging



# Patching







+ ARPA 



**AIxCC**  
AI CYBER CHALLENGE

# PATCHING CRITICAL INFRASTRUCTURE

DARPA's AI Cyber Challenge  
Announced on 08/2023



**AIxCC**  
AI CYBER CHALLENGE

# WHAT IS AIxCC?

- A competition that rewards autonomous systems that *find* and *patch* vulnerabilities in source code — **Cyber Reasoning Systems (CRS)**.
- The challenges are well-known open-source projects.
- The vulnerabilities are realistic or real.
- Patching is worth more than finding.
- Code and data will be released open source.



Preliminary events



Top 7 teams advance



**black hat**

**AUGUST 2023**

**OPEN TRACK AND  
SMALL BUSINESS TRACK  
SUBMISSIONS**



**DEFCON**

**AUGUST 2024**

**SEMIFINAL COMPETITION**

Top 7 teams \$2 million each



**DEFCON**

**AUGUST 2025**

**FINAL COMPETITION**

Winners announced

**1ST: \$4 MILLION**

**2ND: \$3 MILLION**

**3RD: \$1.5 MILLION**

Google

ANTHROPIC


OpenAI

Microsoft

THE  
**LINUX**  
FOUNDATION

**OpenSSF**  
OPEN SOURCE SECURITY FOUNDATION

# What counts for semifinals?



## **Proof-Of-Vulnerability (POV)**


→Input data to reproduce  
vulnerability crash in harness



## **PATCH**


→Unified diff source code  
fix for vulnerabilities

# What counts for finals?



## Proof-Of-Vulnerability (POV)

→ Input data to reproduce vulnerability crash in harness




## SARIF Assessment

→ Structured reporting format for vulnerability details




## PATCH

→ Unified diff source code fix for vulnerabilities



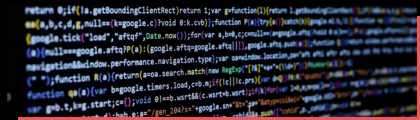
## BUNDLE

→ Grouping of related PoV, patch, and SARIF submissions



## DELTA SCAN

→ Challenge analyzing base code plus applied diff changes



## FULL SCAN

→ Challenge analyzing entire code base

# Open Sourced!



## Artiphishell

Shellphish's cyber reasoning system.

**Team: Shellphish**

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Atlantis

Team Atlanta's cyber reasoning system.

**Team: Team Atlanta**

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## RoboDuck

Theori's cyber reasoning system.

**Team: Theori**

>>> Theori's official AIxCC Landing Page: [GitHub](#)

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Buttercup

Trail of Bits' cyber reasoning system.

**Team: Trail of Bits**

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Fuzzing Brain

All You Need IS A Fuzzing Brain's cyber reasoning system.

**Team: All You Need IS A Fuzzing Brain**

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Bug Buster

42 beyond 6ug's cyber reasoning system.

**Team: 42 beyond 6ug**

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Lacrosse

Lacrosse's cyber reasoning system.

**Team: Lacrosse**

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)

# Open Sourced! But largely unmaintained...

No Activity



## Artiphishell

Shellphish's cyber reasoning system.

**Team:** Shellphish

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Atlantis

Team Atlanta's cyber reasoning system.

**Team:** Team Atlanta

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)

Archived

Archived



## RoboDuck

Theori's cyber reasoning system.

**Team:** Theori

>>> Theori's official AIxCC Landing Page: [GitHub](#)

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Buttercup

Trail of Bits' cyber reasoning system.

**Team:** Trail of Bits

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)

Active

Active



## Fuzzing Brain

All You Need IS A Fuzzing Brain's cyber reasoning system.

**Team:** All You Need IS A Fuzzing Brain

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)



## Bug Buster

42 b3yond 6ug's cyber reasoning system.

**Team:** 42 b3yond 6ug

>>> Semifinal Release: [GitHub](#)

>>> Final Release: [GitHub](#)

No Activity

No Activity



## Lacrosse

Lacrosse's cyber reasoning system.

**Team:** Lacrosse

>>> Semifinal Release: [GitHub](#)


>>> Final Release: [GitHub](#)



OpenSSF Community Day  
NORTH AMERICA 2026

# Running All Finalist CRSs


- ✗ Cloud Lock-In
- ✗ Infrastructure Duplication
- ✗ Monolithic Design



# OSS-CRS: Open CRS Framework




# AIxCC: Score Over Time




# OSS-CRS Features


- Compatible with OSS projects integrated with **OSS-Fuzz**
- Drives both **bug-finding** and **patching** CRSs (and more!)
- Flexible **LLM** endpoint configurations
- **Locally** deployable (run it on your laptop!)
- Resource management and LLM budgeting as first class

# OSS-CRS at a high level






CRSs




Projects




Deployments





Researchers

CRSs




Security Engineers

Projects

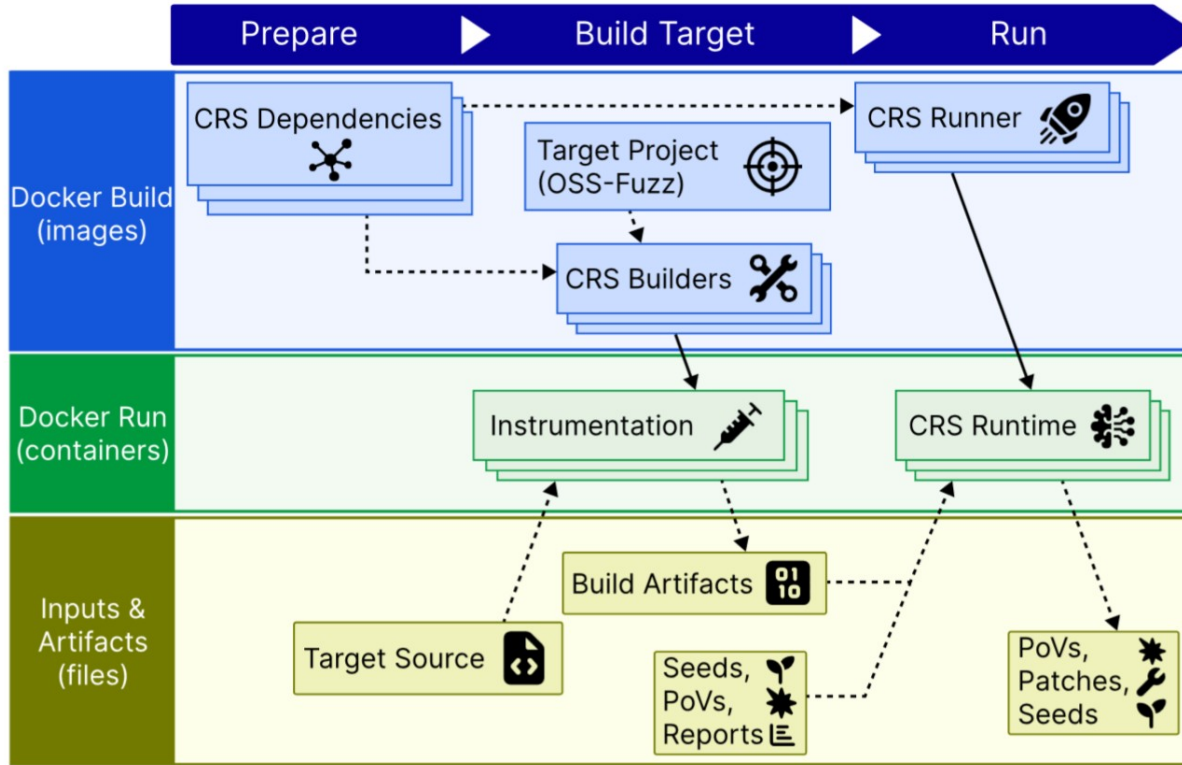


Deployments

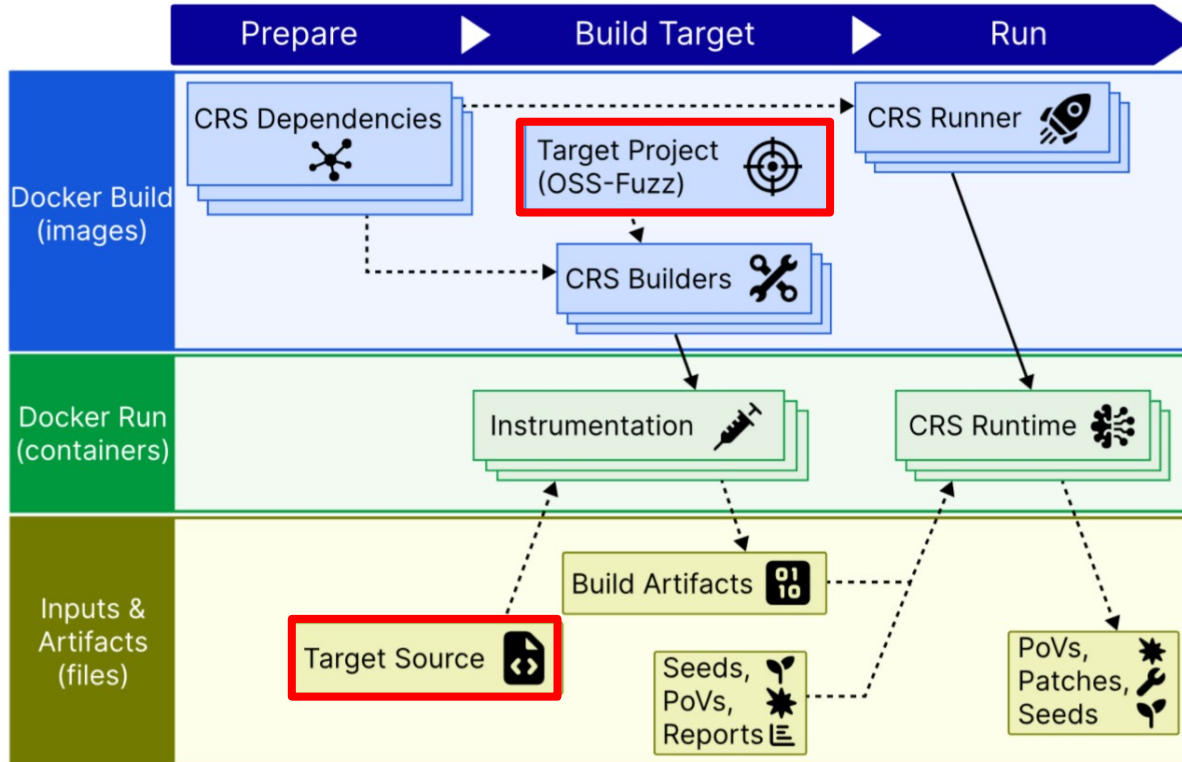
Maintainers



# Stages



# @Project Maintainers



# OSS-Fuzz

- > 1000 OSS projects in the ecosystem
- Standardizes compilation and the execution environment
- Fuzz testing continuously done by Google (Clusterfuzz)
- Threat model defined by “harnesses”



# OSS-Fuzz Format

- Dockerfile – container environment for building the project
- build.sh – script for compiling the project
- Fuzz harnesses – entrypoint that passes data to public APIs
- (Optional) run\_tests.sh – functionality tests

```
1 #include <stdint.h>
2 #include <stdlib.h>
3
4 #include <memory>
5
6 #include <turbojpeg.h>
7
8
9 extern "C" int LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
10     tjhandle jpegDecompressor = tjInitDecompress();
11
12     int width, height, subsamp, colorspace;
13     int res = tjDecompressHeader3(
14         jpegDecompressor, data, size, &width, &height, &subsamp, &colorspace);
15
16     // Bail out if decompressing the headers failed, the width or height is 0,
17     // or the image is too large (avoids slowing down too much). Cast to size_t to
18     // avoid overflows on the multiplication
19     if (res != 0 || width == 0 || height == 0 || ((size_t)width * height > (1024 * 1024))) {
20         tjDestroy(jpegDecompressor);
21         return 0;
22     }
23
24     std::unique_ptr<unsigned char[]> buf(new unsigned char[width * height * 3]);
25     tjDecompress2(
26         jpegDecompressor, data, size, buf.get(), width, 0, height, TJPf_RGB, 0);
27
28     tjDestroy(jpegDecompressor);
29
30     return 0;
31 }
```



# Delta Scan Mode For CI/CD

```
10     int pC2 = (1 * C2_size) + j;  
11     int pD2 = (k * D2_size) + j;  
12     int pA2 = (i * A2_size) + j;  
13     A[pA2] += B[pB3] * C[pC2] * D[pD2];  
14     w[j] += B[pB3] * C[pC2];  
15 }  
16 )  
17 )
```


## DELTA SCAN

Challenge analyzing base  
code plus applied diff changes

Source



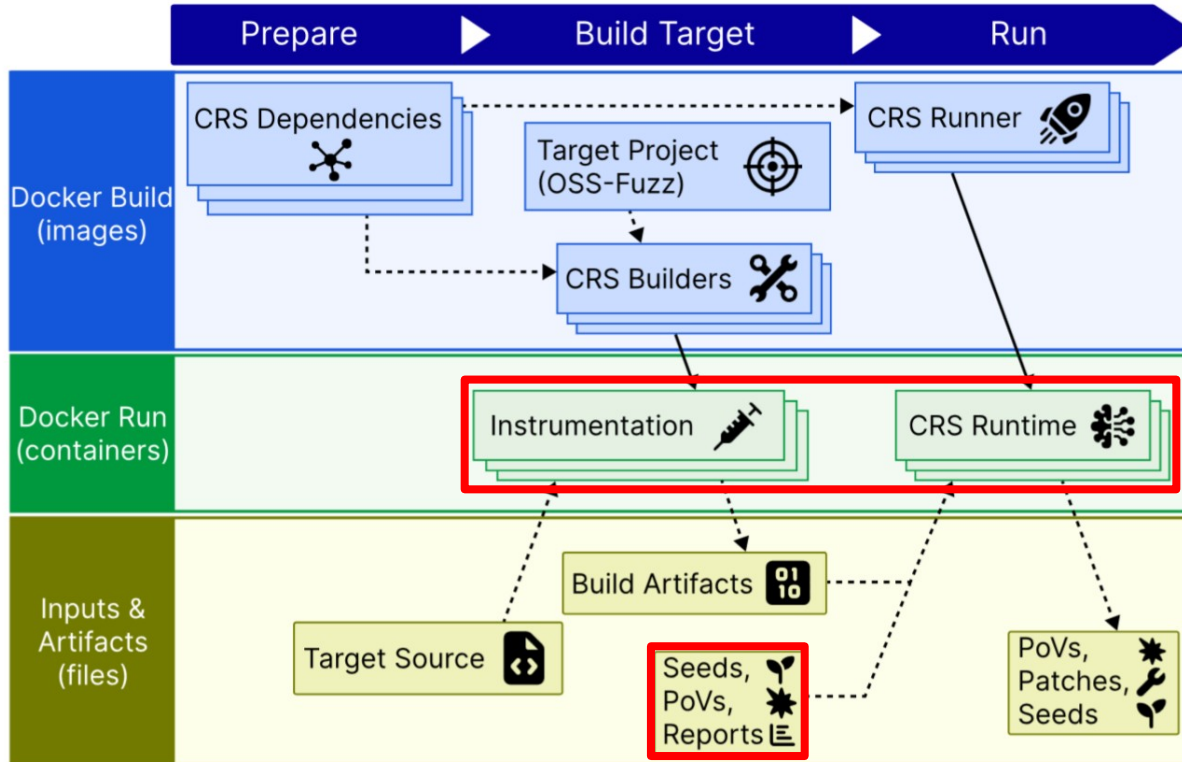

Diff



CRS



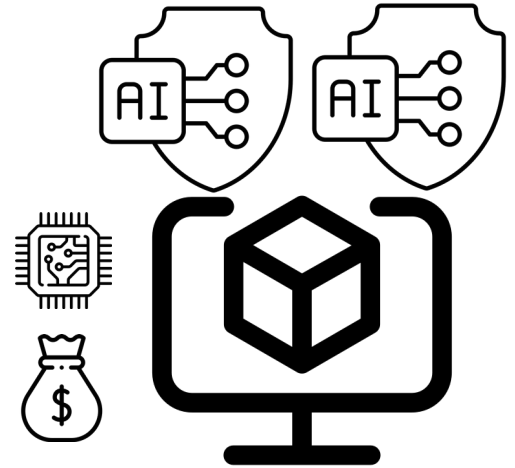
# @Security Engineers



# Using OSS-CRS

The operator of OSS-CRS can decide:

- Which CRSs to run
- How much compute resources to give to each CRS
- How much LLM budget to give to each CRS
- Which models and endpoints to route LLM requests to
- How long to run the CRSs for
- What project (and harness) to run the CRSs against
- Which vulnerability to fix (for patching CRSs)



# Operator Configuration

For a CRS's build and run phase,  
OSS-CRS can apply resource limitations

- CPU set
- Memory
- LLM budget and rate limits

## compose.yaml

```
# `oss-crs-infra` must be required
oss_crs_infra:
  cpuset: "0-3"
  memory: "16G"

crs-libfuzzer:
  cpuset: "8-11"
  memory: "16G"

multilang:
  cpuset: "4-7"
  memory: "16G"
  llm_budget: 100 # 100 dollars

llm_config:
  litellm_config: ./example/multilang/litellm-config.yaml
```

# LLM Configuration

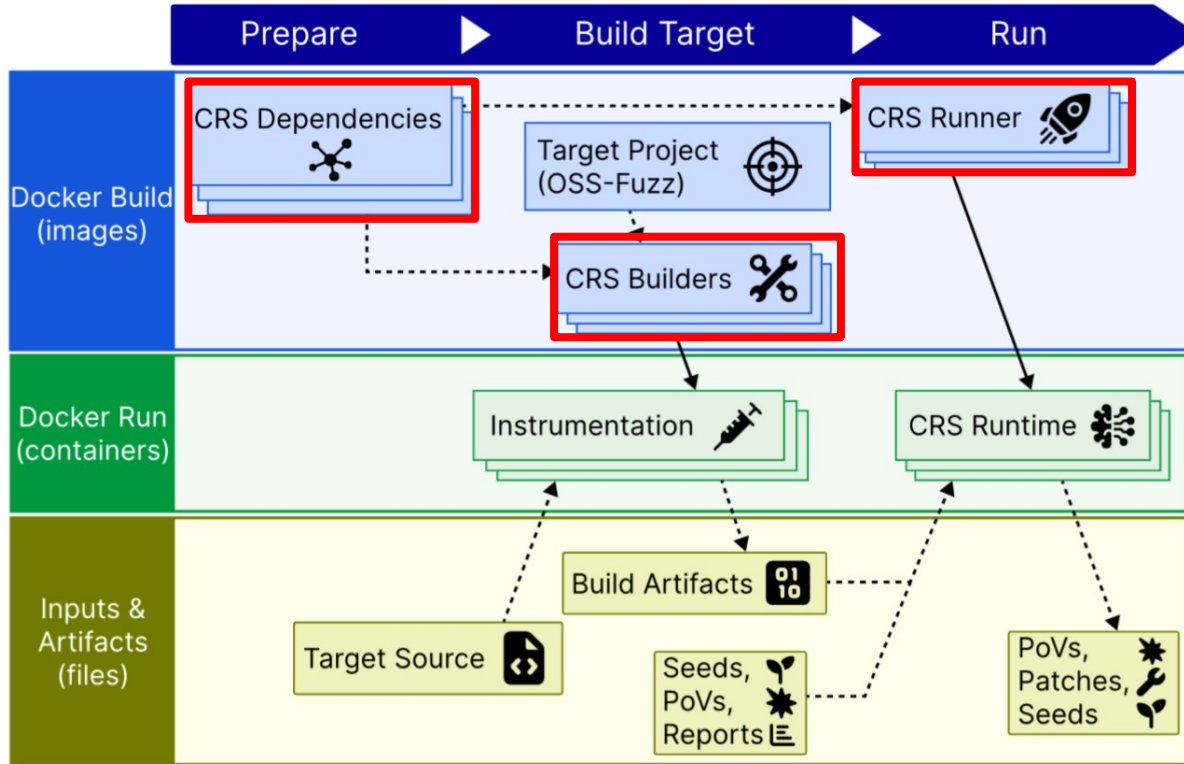

Operators can define a custom configuration for LiteLLM proxy

- Alias LLM model names
- Bring your own API keys

```
model_list:  
- model_name: "claude-opus-4-5-20251101"  
  litellm_params:  
    model: "openai/Qwen/Qwen3-0.6B"  
    api_key: os.environ/VLLM_KEY  
    api_base: https://example.com:8000/v1
```

CRSs should re-route their LLM requests to an OSS-CRS-provided base URL and key

# @Researchers



# CRS Integration Configuration

## Register CRS into OSS-CRS

[ossf/oss-crs: registry/atlantis-java-main.yaml](https://ossf/oss-crs: registry/atlantis-java-main.yaml)

```
name: atlantis-java-main
type:
  - bug-finding
source:
  url: https://github.com/Team-Atlanta/atlantis-java-snapsho
  ref: main
```

## Declare CRS containers and requirements

[Team-Atlanta/atlantis-java-main: oss-crs/crs.yaml](https://Team-Atlanta/atlantis-java-main: oss-crs/crs.yaml)

```
name: atlantis-java-main

target_build_phase:
  # additional containers can be specified
  build:
    dockerfile: oss-crs/builder.Dockerfile

crs_run_phase:
  run:
    dockerfile: oss-crs/runner.Dockerfile

allowed_llms:
- gpt-5
- o4-mini
- o3
- gpt-4.1
```

# libCRS: helpers and interface for CRSs


## Artifact transfer

- Build outputs – syncing build artifacts between stages
- Submitting findings – PoVs, patches
- Fetching auxiliary data – seeds, bug candidates, PoVs, patches

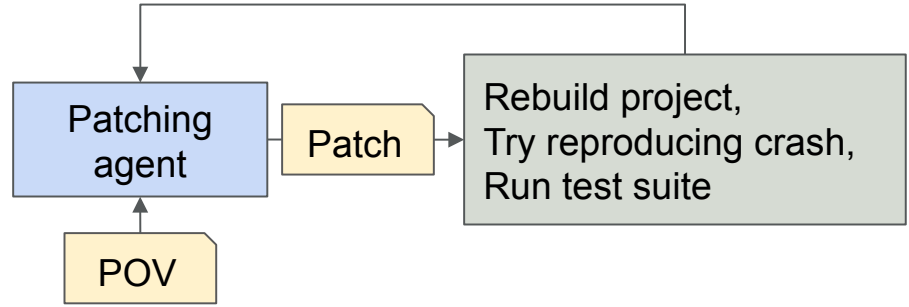
## Helpers for common CRS actions

- Rebuilding the project given a patch – can use custom CRS builder!
- Run PoV to check if buggy behavior reproduces
- Run functionality tests

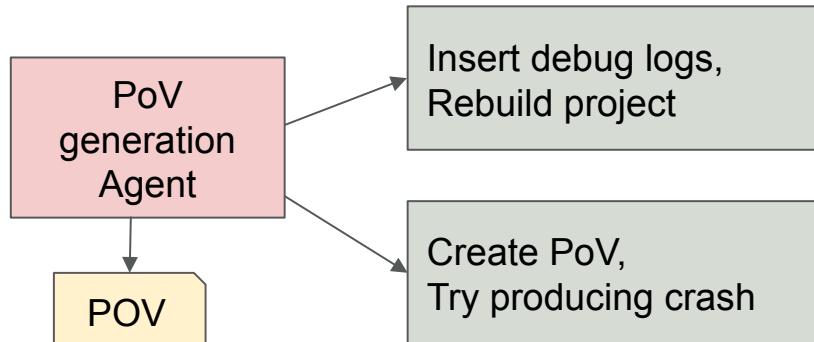
# Common CRS Flows



Seed generation + fuzzer



Patching agent + rebuild



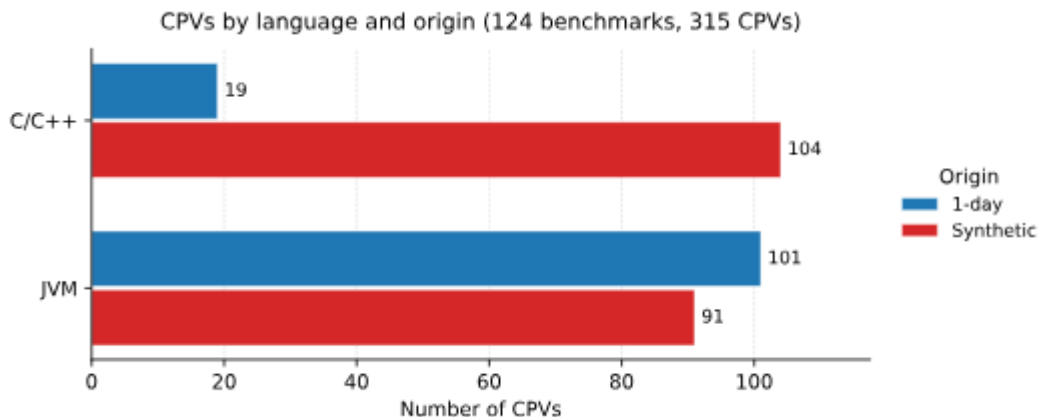
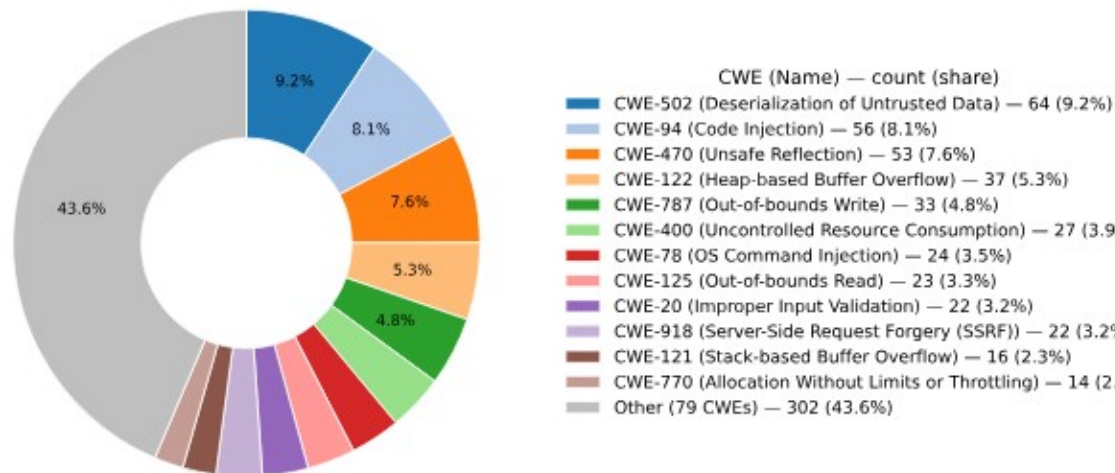
Pure agentic  
PoV generation

# CRSBench

## Benchmarking and testing CRSs

- Use any CRS integrated into OSS-CRS
- Faster evaluation using build snapshots and selective unit testing
- Dataset made from AI Cyber Challenge problems and Team Atlanta internal benchmarks

[oss-crs.openssf.org/crsbench](https://oss-crs.openssf.org/crsbench)



# Putting Everything Together

```
# Build CRS Docker images
```

```
uv run oss-crs prepare \  
  --compose-file example/crs-bug-finding-codex/compose.yaml
```

Compose selects which CRS to run

```
# Build target project image and harnesses
```

```
uv run oss-crs build-target \  
  --compose-file example/crs-bug-finding-codex/compose.yaml \  
  --fuzz-proj-path ~/post/CRSBench/benchmarks/asc-nginx-delta-01
```


“Project” dir with harness and Docker env

```
# Run CRS against a harness (entrypoint)
```

```
uv run oss-crs run \  
  --compose-file example/crs-bug-finding-codex/compose.yaml \  
  --fuzz-proj-path ~/post/CRSBench/benchmarks/asc-nginx-delta-01 \  
  --diff ~/post/CRSBench/benchmarks/asc-nginx-delta-01/.aixcc/ref.diff \  
  --target-harness pov_harness
```

Harness is the entrypoint to public API

Get Involved!



Researchers

CRSs


Security Engineers

Projects



Deployments

Maintainers



# OSS-CRS Links



<https://oss-crs.openssf.org/>



<https://github.com/ossf/oss-crs>



<https://openssf.slack.com/archives/C09FQLYH1RD>

(All of these links can be found on <https://openssf.org/projects/oss-crs/>)



Also check out Team Atlanta's website and blog at <https://team-atlanta.github.io/>



Finally, you can find my contacts on my webpage at <https://achin.ca>



# Thank You!



OpenSSF Community Day  
NORTH AMERICA 2026