

Security Hardening for EPICS/RTEMS

GEDARE BLOOM

UNIVERSITY OF COLORADO COLORADO SPRINGS

EPICS COLLABORATION FALL 2020 MEETING






EPICS is
insecure



EPICS is
insecure?



Industrial Control System Vulnerabilities

Source: <https://ics-cert.kaspersky.com/reports/2020/04/24/threat-landscape-for-industrial-automation-systems-vulnerabilities-identified-in-2019/>. Threat landscape for industrial automation systems, 2019.

EPICS is not protected against sophisticated attackers and unsophisticated interns

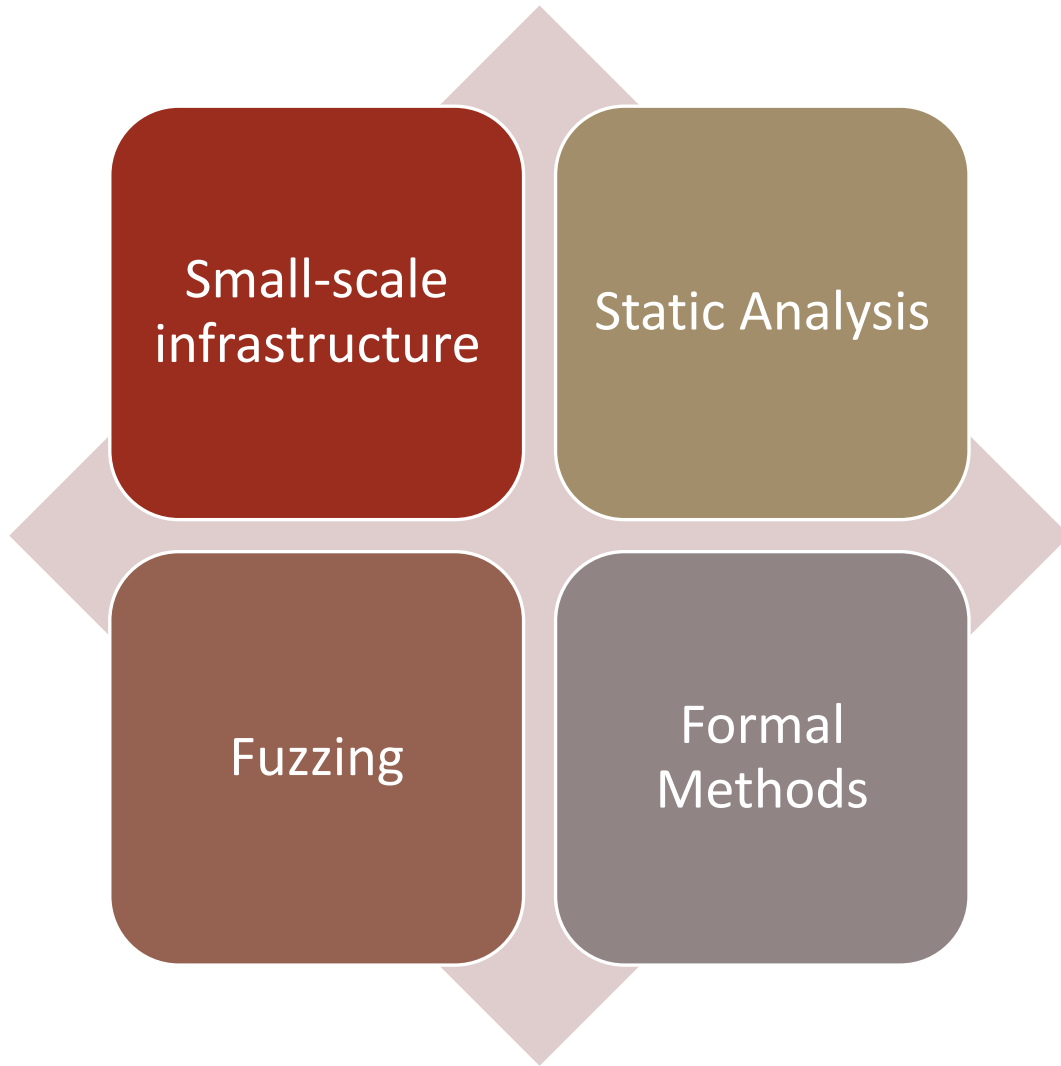
C/C++ will never (!) be secure

May be capable of injecting bugs in EPICS software products

May be capable of launching DOS/DDOS against OPI/PV Gateway/CA networks

May be capable of subverting data quality and provenance of PV data

Goal: provide resiliency for EPICS throughout software development lifecycle



Security Testing

Small-Scale Research Infrastructure

Problem: Want to analyze security problems and solutions with *lean* supplies budget

- Academic lean budget: Supply costs around \$1-2K

Goals:

- Run “realistic” PV/asyn workloads
- Evaluate runtime/security performance of modifications to EPICS (IOCs: RTEMS, Linux)
- Student setup and maintenance

Real-time Analytics with Open Powerlink


Investigating system delay, network delay, actuation delay

- System delay: background tasks, user / kernel space
- Propagation delay: interference in wire
- Actuation delay: servo response time; screw, motor

Real-time analytics in an EPICS environment

- IOCs → OPI → On-Site Private Cloud (OSPC)
- Extension of our baremetal PowerLink stack
 - Video → PowerLink management node polling controlled node

Prajjwal Dangal, Gedare Bloom, Towards Industrial Security Through Real-time Analytics, in 2020 IEEE 23rd International Symposium on Real-Time Distributed Computing (ISORC), pp. 156-157, 2020.





EPICS Simulation package with utilities for data acquisition, exporting, noise setting [2]



Ophyd provides EpicsMotors, Signal classes backed by PVs



Databroker, event driven



No physical test-bed required

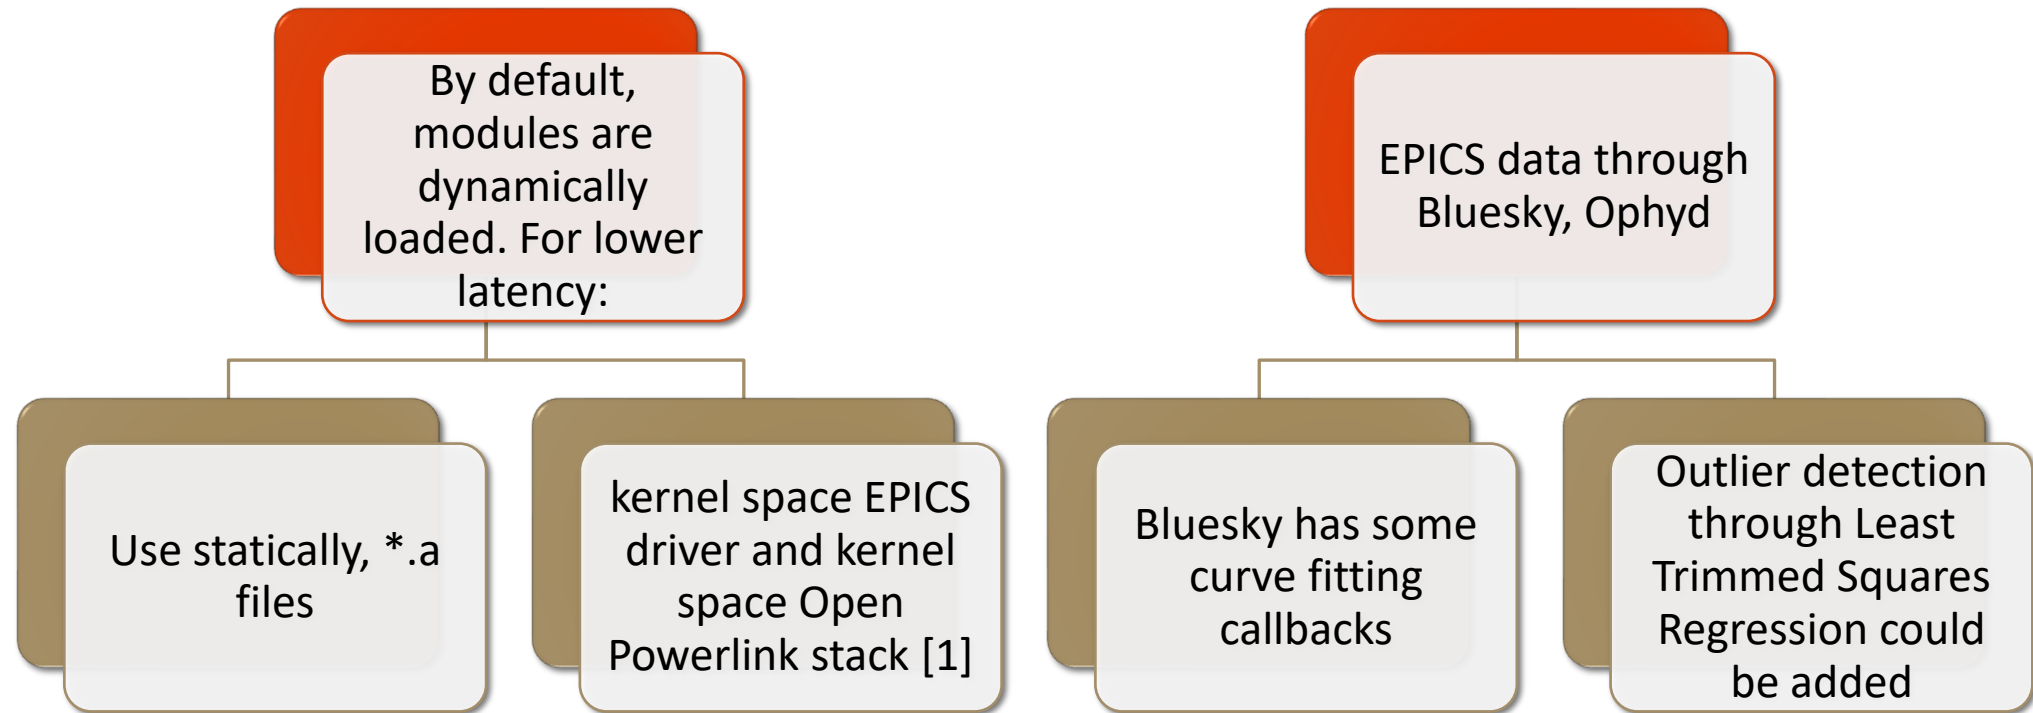


Easy to reconfigure, update

Bluesky testbed-free infrastructure

<https://blueskyproject.io/>

Current Approach with PowerLink



Bluesky, “bluesky/tutorials,” <https://github.com/bluesky/tutorials>
beamlineconfiguration.py is the king 😊

Next Steps for Research Infrastructure



Using Javascript for Bluesky's data model backend as it uses documents



Port data analysis packages related to linear, logistic regression, etc.



Working with EPICS and RTEMS for real-time analytics



Decorate, modify Bluesky code through Numba, and Numpy

Static Analysis

- ✓ Analyze the software without executing
- x Cannot find runtime and configuration errors
- x Cannot test EPICS network

First approach: Coverity

Mature product

Open-source friendly

Kind of

Hooks with Travis-ci

Pain to extract defects

Investigating with RTEMS

Second approach: clang-analyzer

Active open-source project

Lot of other analyzers – also trying

Static Value flow (SVF) analyzer

Easy to work with

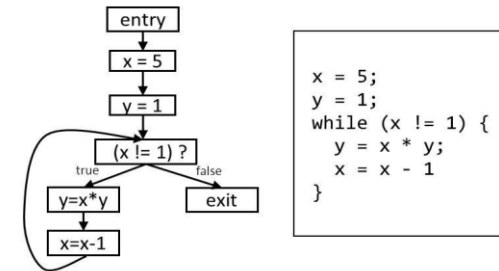
Defect report is like compiler

Must compile with clang

Used it with RTEMS

Hopeful to apply to EPICS soon!

Control-Flow Graphs



Security Fuzzing

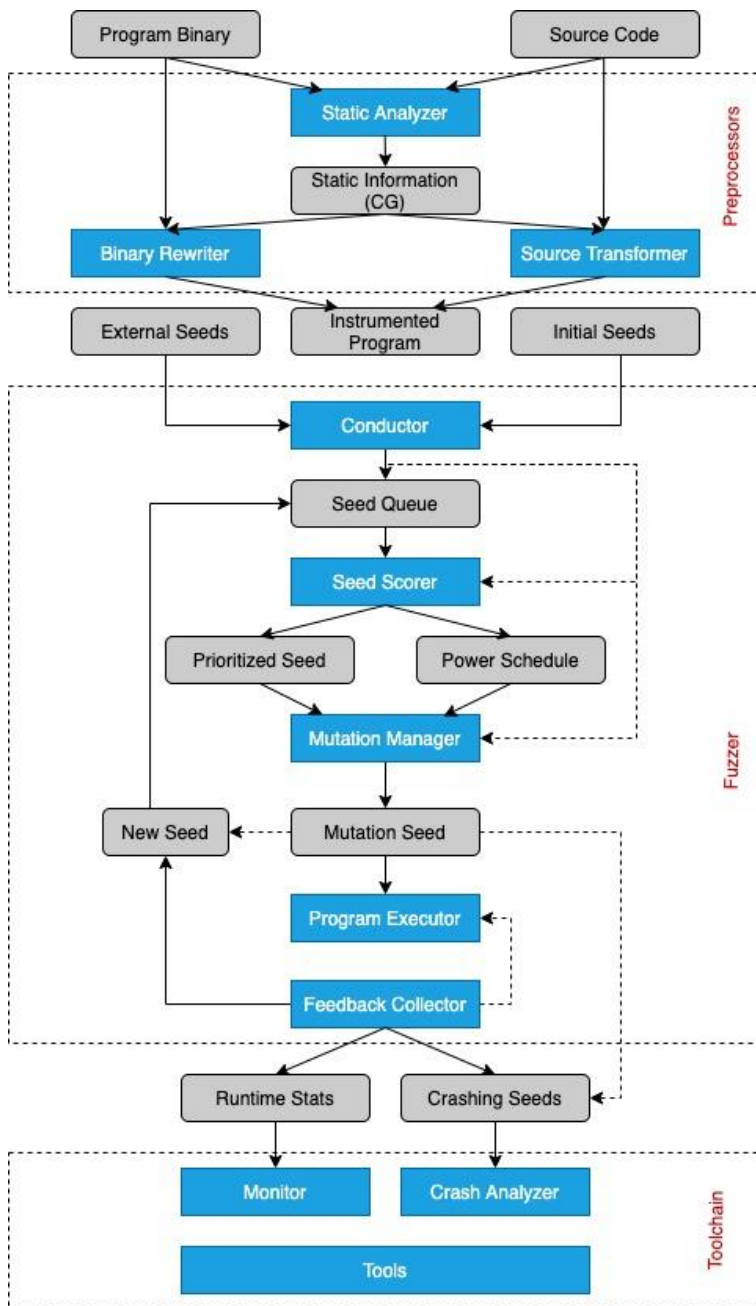
- ✓ Detect security loopholes and vulnerabilities on EPICS System
- ✓ Can detect zero-day vulnerabilities
- x Requires significant time to detect Bugs and vulnerabilities

Current state of the art is American Fuzzy Lop (AFL) fuzzer using mutation-based fuzzing

- As opposed to generation-based fuzzing

Our approach aims to improve AFL and apply to EPICS modules and CA

- Extension of AFL, AFL-NET, MOPT, UniFuzz and MultiFuzz to work for CA



Improve Fuzzing with Static Analysis

```

uchenna@uchenna-Precision-3630-Tower: ~/EPICS_FUZZY/epics-base/fuzzer_file 101x25

american fuzzy lop 2.52b (softIoc)

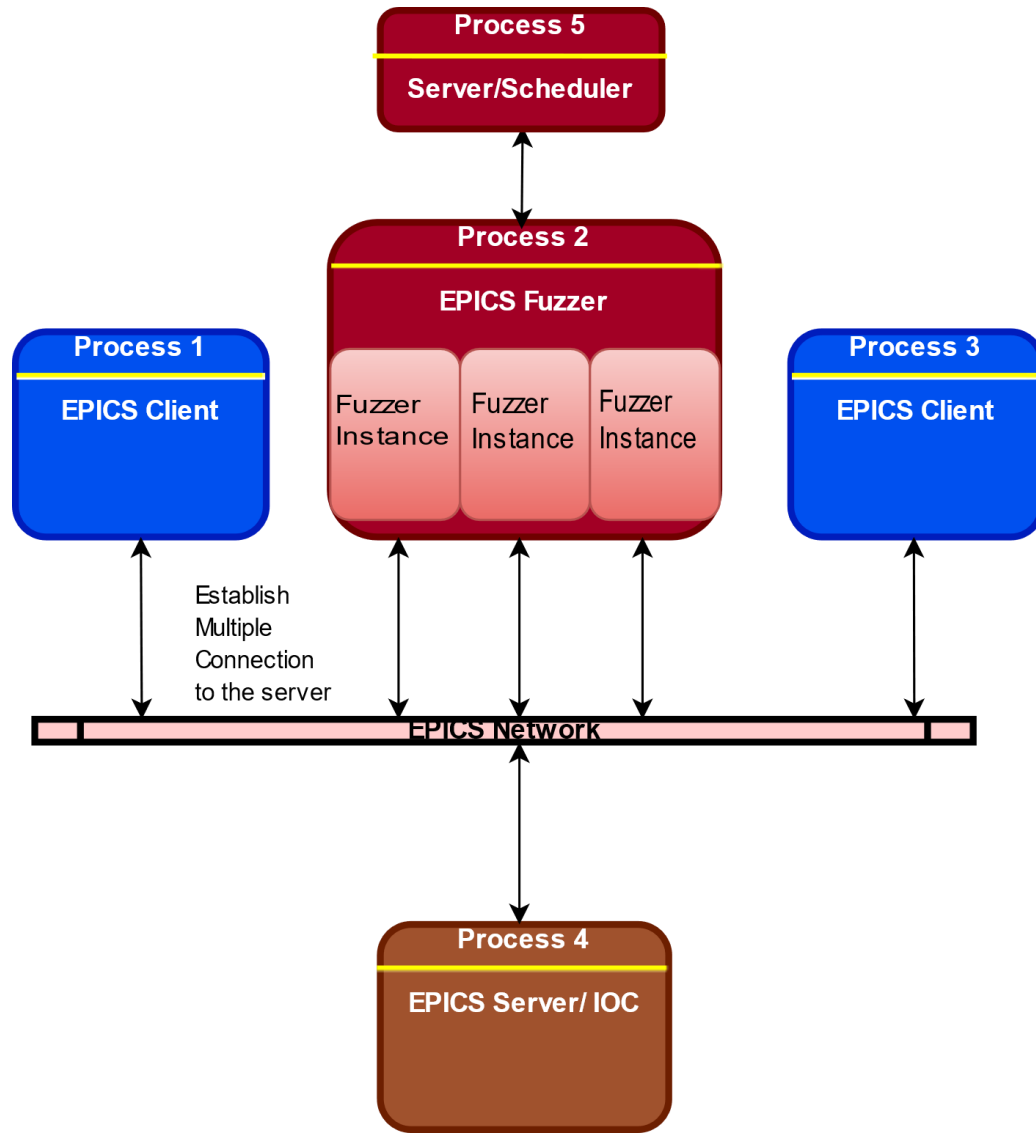
process timing
  run time : 0 days, 17 hrs, 51 min, 54 sec
  last new path : 0 days, 0 hrs, 5 min, 48 sec
  last uniq crash : 0 days, 0 hrs, 3 min, 5 sec
  last uniq hang : 0 days, 1 hrs, 28 min, 56 sec
cycle progress
  now processing : 405 (84.38%)
  paths timed out : 0 (0.00%)
stage progress
  now trying : interest 32/8
  stage execs : 549/750 (73.20%)
  total execs : 515k
  exec speed : 7.92/sec (zzzz...)
fuzzing strategy yields
  bit flips : 60/20.8k, 33/20.7k, 19/20.5k
  byte flips : 0/2603, 1/2490, 1/2265
  arithmetics : 64/145k, 2/27.2k, 1/5785
  known ints : 3/12.6k, 4/65.9k, 8/96.9k
  dictionary : 0/0, 0/0, 39/7929
  havoc : 261/79.8k, 0/0
  trim : 6.81%/605, 0.00%

overall results
  cycles done : 0
  total paths : 480
  uniq crashes : 23
  uniq hangs : 16
map coverage
  map density : 2.86% / 3.64%
  count coverage : 1.92 bits/tuple
findings in depth
  favored paths : 79 (16.46%)
  new edges on : 147 (30.62%)
  total crashes : 105 (23 unique)
  total tmouts : 174 (61 unique)
path geometry
  levels : 9
  pending : 368
  pend fav : 32
  own finds : 473
  imported : n/a
  stability : 96.31%

[cpu004: 58%]
  
```

AFL running on softloc

Current Research: Fuzzing Channel Access



The Fuzzer will have 1 scheduler/ server with N number of fuzzer (clients)

Each N clients can have M instances of Fuzzer (modified version of AFL)

Overload the network to investigate network vulnerabilities

Compare with existing network mutation fuzzer

Challenge: How to store multiple connections and synchronize paths covered and vulnerabilities found between fuzzers?

Formal Methods

THEOREM PROVING

Proof assistance to confirm/deny theorems formulated about system

- Manually labor intensive to formulate theorems so prover can reason over them

Size still a problem. Current approaches handle ~10k's lines of code

Strong expressivity in higher order logic

A few stable tools available

- Coq, HOL4, Isabella, PVS

MODEL CHECKING

Automatically check (verify) properties over an abstraction of system

- Brute force/heuristic approach without manual intervention

Constrained by size – state space explosion


Limited expressivity based on first order logic

- Usually enough for **safety** properties

Widespread tools available

- SPIN, TLA+, UPAAL, PRISM, NuSMV, ...

Model checking in a nutshell



4

Our Approach

Use model checking to find bugs and concurrency problems

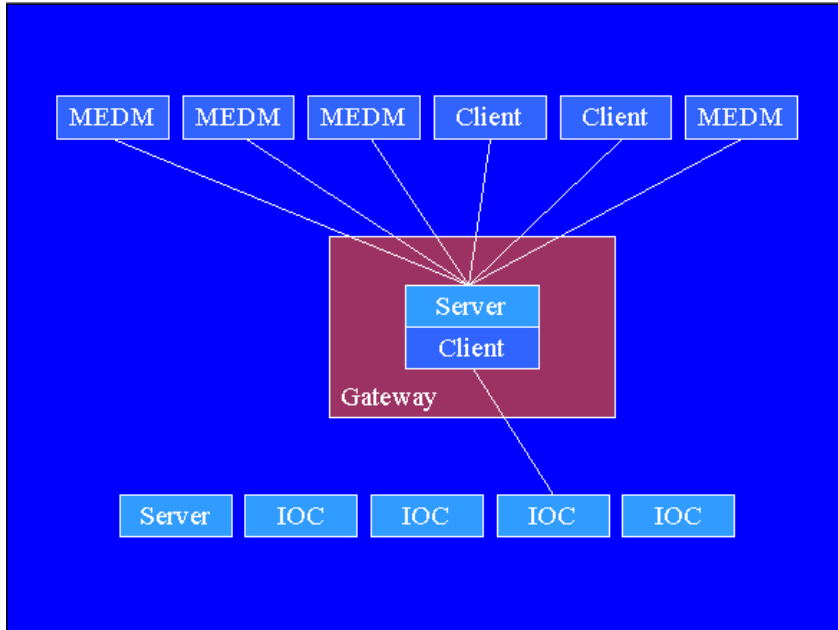
- Analyze PV Gateway using SPIN model checker

Why PV Gateway?

Widely used in EPICS community

Interesting (to us) Threat Model:

- Possibility of Dolev-Yao adversary
 - Network-level control (eavesdrop, add/drop messages)
- Written in C and all its glory
- Uses UDP → potential vulnerability to stateless attacks and UDP amplification attacks



Next Steps in Formal Methods

Finish initial modeling of PV Gateway

- Formulate some interesting properties to check if:
 - MEDM(client) can be restricted in their access to the PV
 - Gateway internal PV used for diagnostic purposes cannot be used as a backdoor?
 - PV_obj is destroyed after 120s unless it is connected?
 - PV_obj in inactive state for more than 7200s.
 - Transition between states are correct. For example, PV_obj in dead state should not be active in less than 120s.

Begin to model state space of CA/PVA protocols

Conclusion

Much work remaining to be done in these areas

Some next steps

- Security in open-source SDLC with signed commits, releases
 - Need to discuss with EPICS core and RTEMS maintainers
- Porting secure communications (dropbear ssh)
- Memory protection – Recurrent topic in RTEMS
 - Looking at how to create a model that fits across IOCs
- Secure Boot
 - ditto