

1. Introduction & Background

- **System call:** Talk to hardware through kernel (Figure 1)
- **Sandboxing:** Restrict system calls to minimal required set (Figure 2)
- **Problem:** Which calls to block and which ones to allow?
- **Solution:** Analyse applications, find out which calls are needed
- **Gap:** Static vs Dynamic analysis & Single vs Multi phase model (Figure 3)

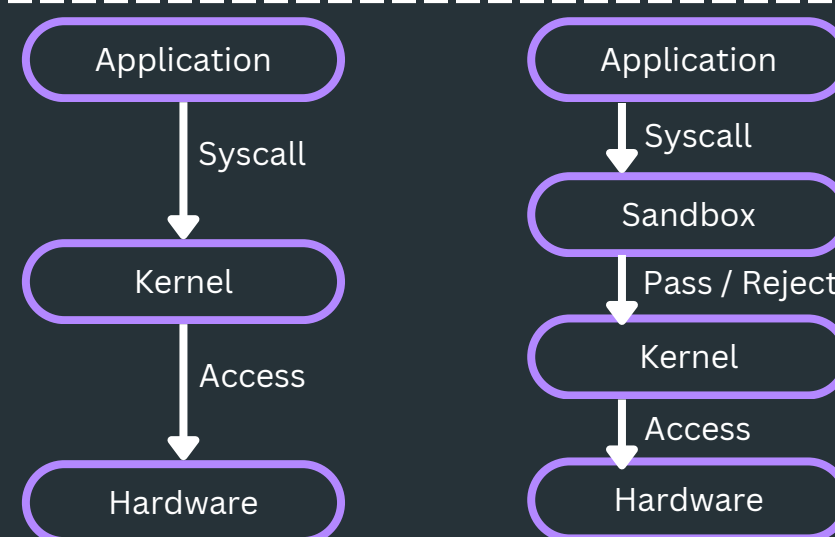



Figure 1: Regular syscall flow

Figure 2: Sandboxed flow

Single phase model



Multi phase model




Figure 3: Single phase & multi phase model

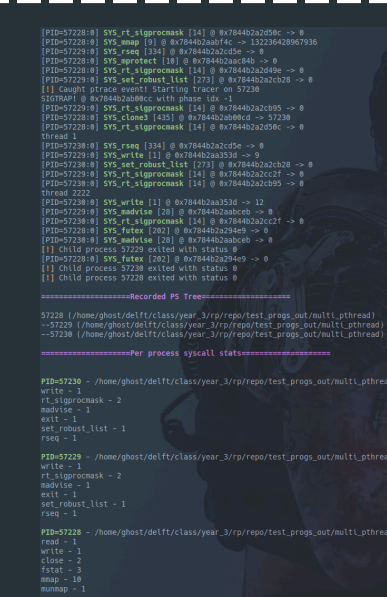


Figure 4: The dynamic tracer

2. Research questions & Contributions

1. How can dynamic analysis method be used to identify the used system calls?
2. What is the runtime and accuracy of single phase model static analysis tools?
3. What is the runtime and accuracy of multi phase model static analysis tools?

3. Dynamic tracer

- Gather list of system calls used, using ptrace
- Traces multiple processes & threads
- Records process structure and system calls (Figure 4)
- Supports multi phase tracing
- Requires exploration of many program states

4. Experiment setup

Test programs: ls (v8.28), sqlite3 (v3.22), redis (v4.0.9)


Environment: Ubuntu 18.04 Docker container

Analysis tools: Chestnut [1], Confine [4], temporal-specialization [3], sysfilter [2]


Measurements: Runtime, Accuracy



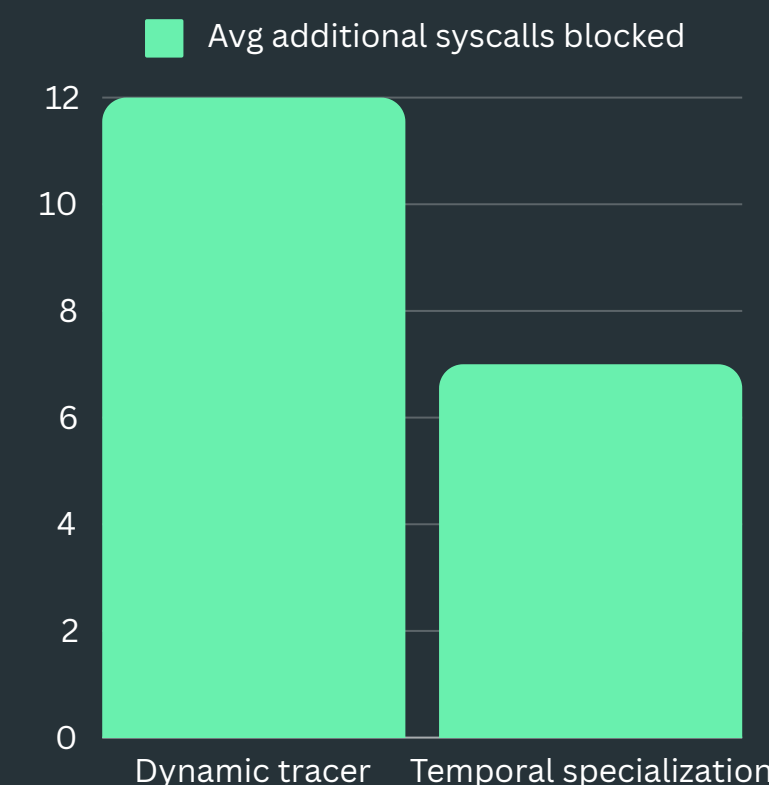
5. Results & Discussion




Dynamic tracer blocks **14%** more syscalls, than Sourecaalyzer on average



Dynamic tracer is **4.5x** faster than Confine on average



Multi phase analysis blocks **11** more syscalls during the second phase than single phase analysis on average



Sourcealyzer is hardest to set up due to little setup instructions and a lot of compiling

6. Conclusion

Dynamic analysis

- +** Allows for custom usage profile
- +** Fast analysis times
- Requires extensive test cases
- i** Second step analysis to reduce amount of incorrectly blocked system calls

Static analysis

- +** Good amount of blocked system calls, with slight underestimation
- Slower analysis time as programs scale
- Difficult setup and maintainability
- i** Pre-computing CFG and parallelism

Multi phase model

- +** Better security through more fine-grained filters
- Need to determine transition point manually
- i** Instruction level transition points
- i** Multiple transition points

Possible future work

- i** Investigate if programs lose any functions after applying the filters
- i** Expand the set of tested programs

References

- [1] Claudio Canella, Mario Werner, Daniel Gruss, and Michael Schwarz. Automating seccomp filter generation for linux applications. In Proceedings of the 2021 on Cloud Computing Security Workshop, CCSW '21, page 139–151, New York, NY, USA, 2021. Association for Computing Machinery.
- [2] Nicholas DeMarinis, Kent Williams-King, Di Jin, Rodrigo Fonseca, and Vasileios P. Kemerlis. sysfilter: Automated system call filtering for commodity software. In International Symposium on Recent Advances in Intrusion Detection, 2020.
- [3] Seyedhamed Ghavamnia, Tapti Palit, Shachee Mishra, and Michalis Polychronakis. Temporal system call specialization for attack surface reduction. In Proceedings of the 29th USENIX Conference on Security Symposium, SEC'20, USA, 2020. USENIX Association.
- [4] Seyedhamed Ghavamnia, Tapti Palit, Azzedine Benameur, and Michalis Polychronakis. Confine: Automated system call policy generation for container attack surface reduction. In International Symposium on Recent Advances in Intrusion Detection, 2020