

Tarnhelm: Isolated, transparent and confidential execution of arbitrary code in ARM's TrustZone



*Davide Quarta, Michele Ianni, **Aravind Machiry**, Yanick Fratantonio, Eric Gustafson,
Davide Balzarotti, Martina Lindorfer, Giovanni Vigna, Christopher Kruegel*



Overview

1. Introduction
2. Trusted Execution Environments
3. Design Goals
4. Approach
5. Implementation
6. Security Evaluation
7. Performance Evaluation
8. Conclusion

Introduction

Introduction

Applications running on a commodity operating system are usually deployed in an untrusted environment.

The user has full access to any of the application's assets, *including its code*.

Introduction

In the absence of architectural support to protect an application's code from unauthorized access, thus avoiding intellectual property loss and piracy of paid content, developers have to rely on:

- Code obfuscation
- Anti-tampering and Anti-debugging techniques
- Different distribution strategies (e.g., in-app purchases)

**“All intellectual property protection technologies will be cracked at some point
- it’s just a matter of time”
- *Microsoft***

Can we achieve **Code Confidentiality**
using **Trusted Execution Environments**?

Introduction

- TEEs operate on a higher level of privilege, they are only designed to execute trusted code signed by device vendors
- TEEs are resource-constrained and not designed to execute full-fledged applications


We address these challenges in **Tarnhelm**, which *transparently* executes individual code components in TrustZone and guarantees *code confidentiality* through isolation, without sacrificing overall system security.

Trusted Execution Environments

Trusted Execution Environment (TEE)

- Hardware-isolated execution environment (e.g., ARM TrustZone)
 - Non-secure world
 - Untrusted OS and untrusted applications (UAs) (e.g., Android and apps)
 - Secure world
 - Higher privilege, can access *everything*
 - Trusted OS and trusted applications (TAs)

ARM TrustZone



Limitations of Existing TEEs

Developers must

- manually partition an application's code into a secure and non-secure part;
- define interfaces between the two parts;
- modify the secure code part to be compatible with the TEE.


Design Goals

Design Goals

- Code confidentiality
- Transparent forwarding
- Transparent integration
- Limited attack surface
- Minimal overhead

Approach


Deployment




Code Partitioning

```
1  #include<stdio.h>
2  int curr_idx = 0;
3  + #define __tarnhelm __attribute__((section(".invisible")))
4  + __tarnhelm void* get_processed_data(struct object *data){
5
6  - void* get_processed_data(struct object *data){
7
8     increment_counter(data);
9     // use data to perform some computation
10    return data;
11 }
12 void increment_counter(struct object *data){
13     if(data != NULL){
14         data->counter += curr_idx;
15         curr_idx++;
16     }
17 }
18 int main(){
19     struct object curr_data;
20     ...
21     get_processed_data(curr_data);
22     ...
23 }
```


Secure Code Retrieval and Loading




Memory Management



System Call Forwarding



Transparent World Switch




Implementation

Implementation

We implemented Tarnhelm based on the default OP-TEE 2.3.0 32-bit QEMU configuration. We added:

- 3.11K lines of code (LOC) to the TCB
- 1,415 LOC to the OP-TEE OS
- 566 LOC to the Linux abort handler and include files
- 1,129 LOC to the OP-TEE Linux driver

Transparent Execution



Control-Flow Integrity

↓From/To→		Untrusted OS	Trusted OS
Untrusted OS	ret	N/A	Verify and pop the return address from the shadow stack
	call	N/A	Verify function entry point and push return address on the shadow stack
Trusted OS	ret	Pop shadow stack	Verify return location to be valid
	call	Push return address on the shadow stack	Verify function entry point for indirect calls

Security Evaluation

Attacks on Code Confidentiality

- Instruction inference attacks
- Control-flow redirection attacks
- Data-only attacks
- Iago attacks
- Blind ROP
- Vulnerabilities in the invisible code
- Compromised TA
- Emulated TEE

Performance Evaluation

Performance Evaluation

We evaluated Tarnhelm on QEMU emulating an ARMv7 Cortex-A15 with soft-mmU, running on an Intel Core 8-core i7-930 CPU (2.80GHz) desktop machine with 12GB of memory.


Microbenchmark of Tarnhelm's Individual Components

Component	Time
Invisible code initialization	0.316s
Invisible code cleanup	0.44ms
System call forwarding	116.88 μ s
Data mapping (secure world)	71 μ s
Data mapping (normal world)	231.337 μ s
IW-CFI indirect call (trusted OS)	0.111 μ s
IW-CFI return (trusted OS)	19.431 μ s



Overhead of the Transparent World Switch

Direction	w/ DM+IWCFI	w/ DM fwd	w/o DM fwd
SW $\xrightarrow{\text{call}}$ NW $\xrightarrow{\text{ret}}$ SW	495.529 μ s	494.539 μ s	152.093 μ s
NW $\xrightarrow{\text{call}}$ SW $\xrightarrow{\text{ret}}$ NW	505.348 μ s	497.549 μ s	151.298 μ s
SW $\xrightarrow{\text{id-call}}$ NW $\xrightarrow{\text{ret}}$ SW	514.903 μ s	N/A	N/A

LMBench Results



Macro Experiment with a Real-World Game



Conclusion

Conclusion

- Tarnhelm, an approach that offers a new powerful primitive: code confidentiality
- Transparent execution of parts of an unmodified application in different isolated execution environments
- Limited additions to the TCB
- Resiliency of Tarnhelm against potential attacks
- Reasonable performance overhead
- Open source, available at <https://github.com/ucsb-seclab/invisible-code>

Questions?
