

Compartmentalisation Workshop

Ben Laurie, David Chisnall

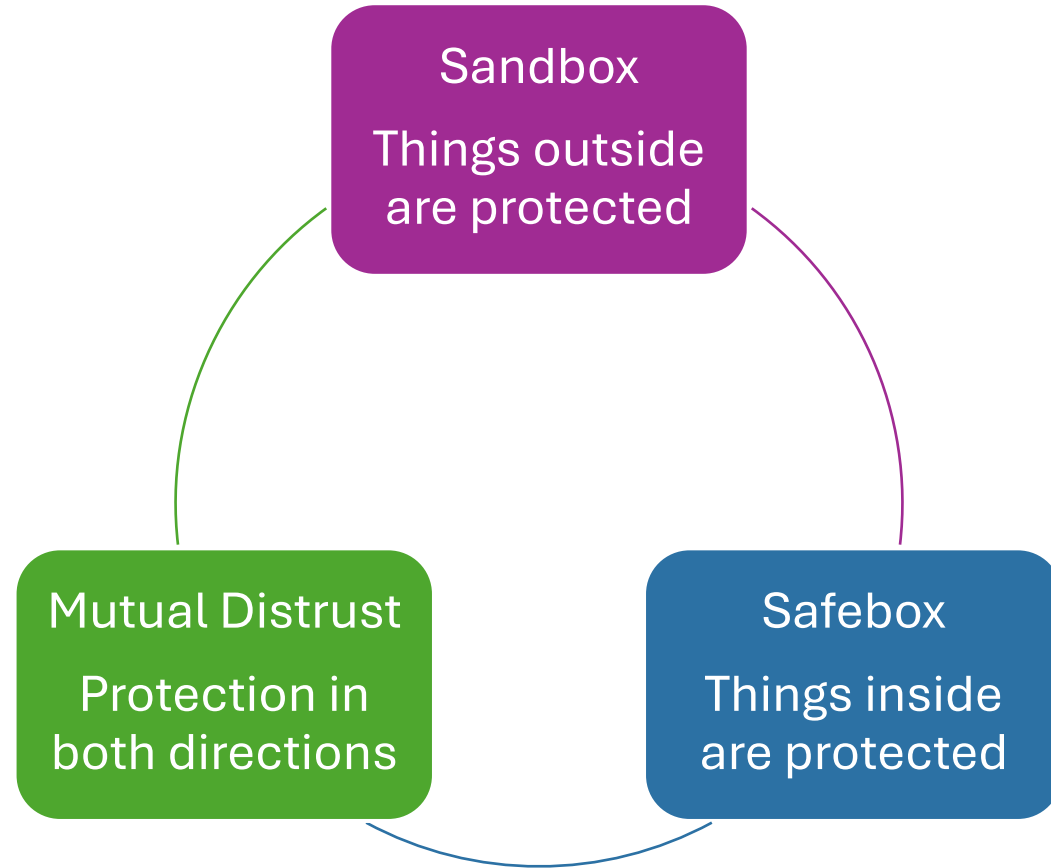


Reducing blast radius

- What breaks if this crashes?
- What breaks if this is compromised?
- What is leaked in either case?



Trust models for compartments

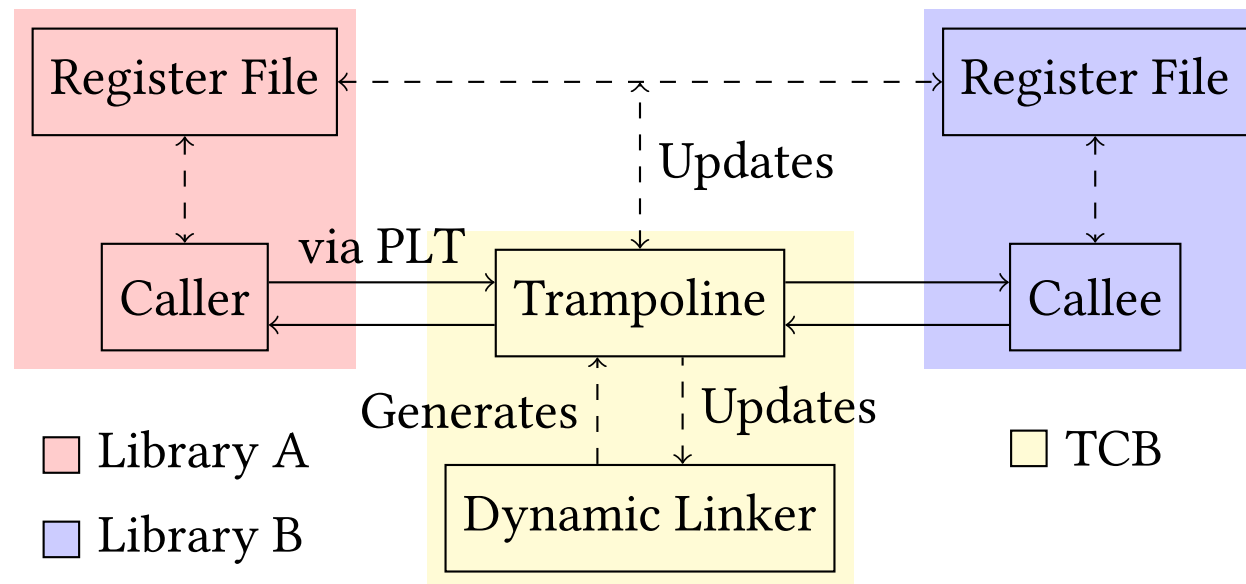


The principle of least privilege

No compartment has rights that it doesn't need

Library-based compartmentalisation

- Each dynamic library is a compartment
- Dynamic linker sets up *trampolines* between cross-compartment calls
 - Clean up unused registers
 - Switch the execution stack
 - Etc.



Demo: Protecting against stack corruption

A `pure_computation` function is supplied by some untrusted third-party library.

```
int pure_computation();
```

This function is supposed to have no side effects, but how do we guarantee that?

The concrete implementation of C implicitly confers many more powerful capabilities to the untrusted function. For example, it could overwrite the stack frame of its caller...

```
cheribsd-d 1 sh* 2 sh-
#include <stdio.h>
#include <string.h>

#define MSG "Hello, world!"

void f();

int main() {
    char buf[14];
    memcpy(buf, MSG, sizeof(MSG));
    printf("%#p\n", buf);
    puts(buf);
    f();
    puts(buf);
    return 0;
}

#include <stdio.h>

void f() {
    char *p;
    asm ("mov    %0, csp" : "=C" (p));
    printf("%#p\n", p);
    p[124] = 'M';
}
~
~
~
~
~
~
~
```

normal:

```
cc main.c f.c -o normal
```

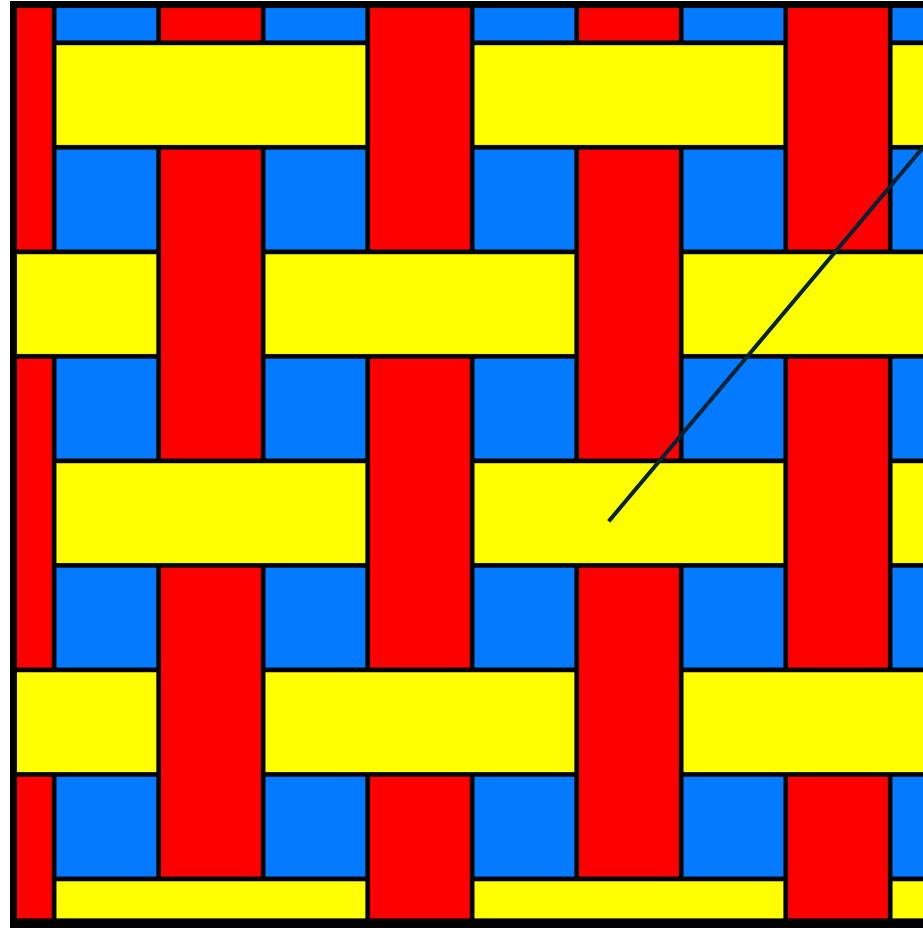
compart:

```
cc f.c -shared -o f.so
cc main.c f.so -o compart -Wl,--dynamic-linker=/libexec/ld-elf-c18n.so.1 -Wl,-rpath=.
```

```
dpgao@cheribsd-dg612:~/src/tap_demo $ ls -lha
total 28
drwxr-xr-x  2 dpgao dpgao  512B Nov  8 13:41 .
drwxr-xr-x 12 dpgao dpgao  512B Nov  5 13:20 ..
-rw-r--r--  1 dpgao dpgao  154B Sep 15 16:45 Makefile
-rw-r--r--  1 dpgao dpgao  112B Sep 15 16:45 f.c
-rwxr-xr-x  1 dpgao dpgao  5.2K Nov  8 13:35 f.so
-rw-r--r--  1 dpgao dpgao  206B Sep 15 16:45 main.c
dpgao@cheribsd-dg612:~/src/tap_demo $
dpgao@cheribsd-dg612:~/src/tap_demo $ make normal
```

CHERIoT has two-dimensional isolation

Compartments
own code and
globals



Running code can access
the current compartment's
state, on behalf of the
current thread

Threads are scheduled independently and call through compartments

Compartmentalisation with CHERIoT

- Compartments are invoked as function calls / returns
- Safe return is guaranteed by a *trusted stack*
- Callees can access explicit arguments
- Callees cannot access any other caller state



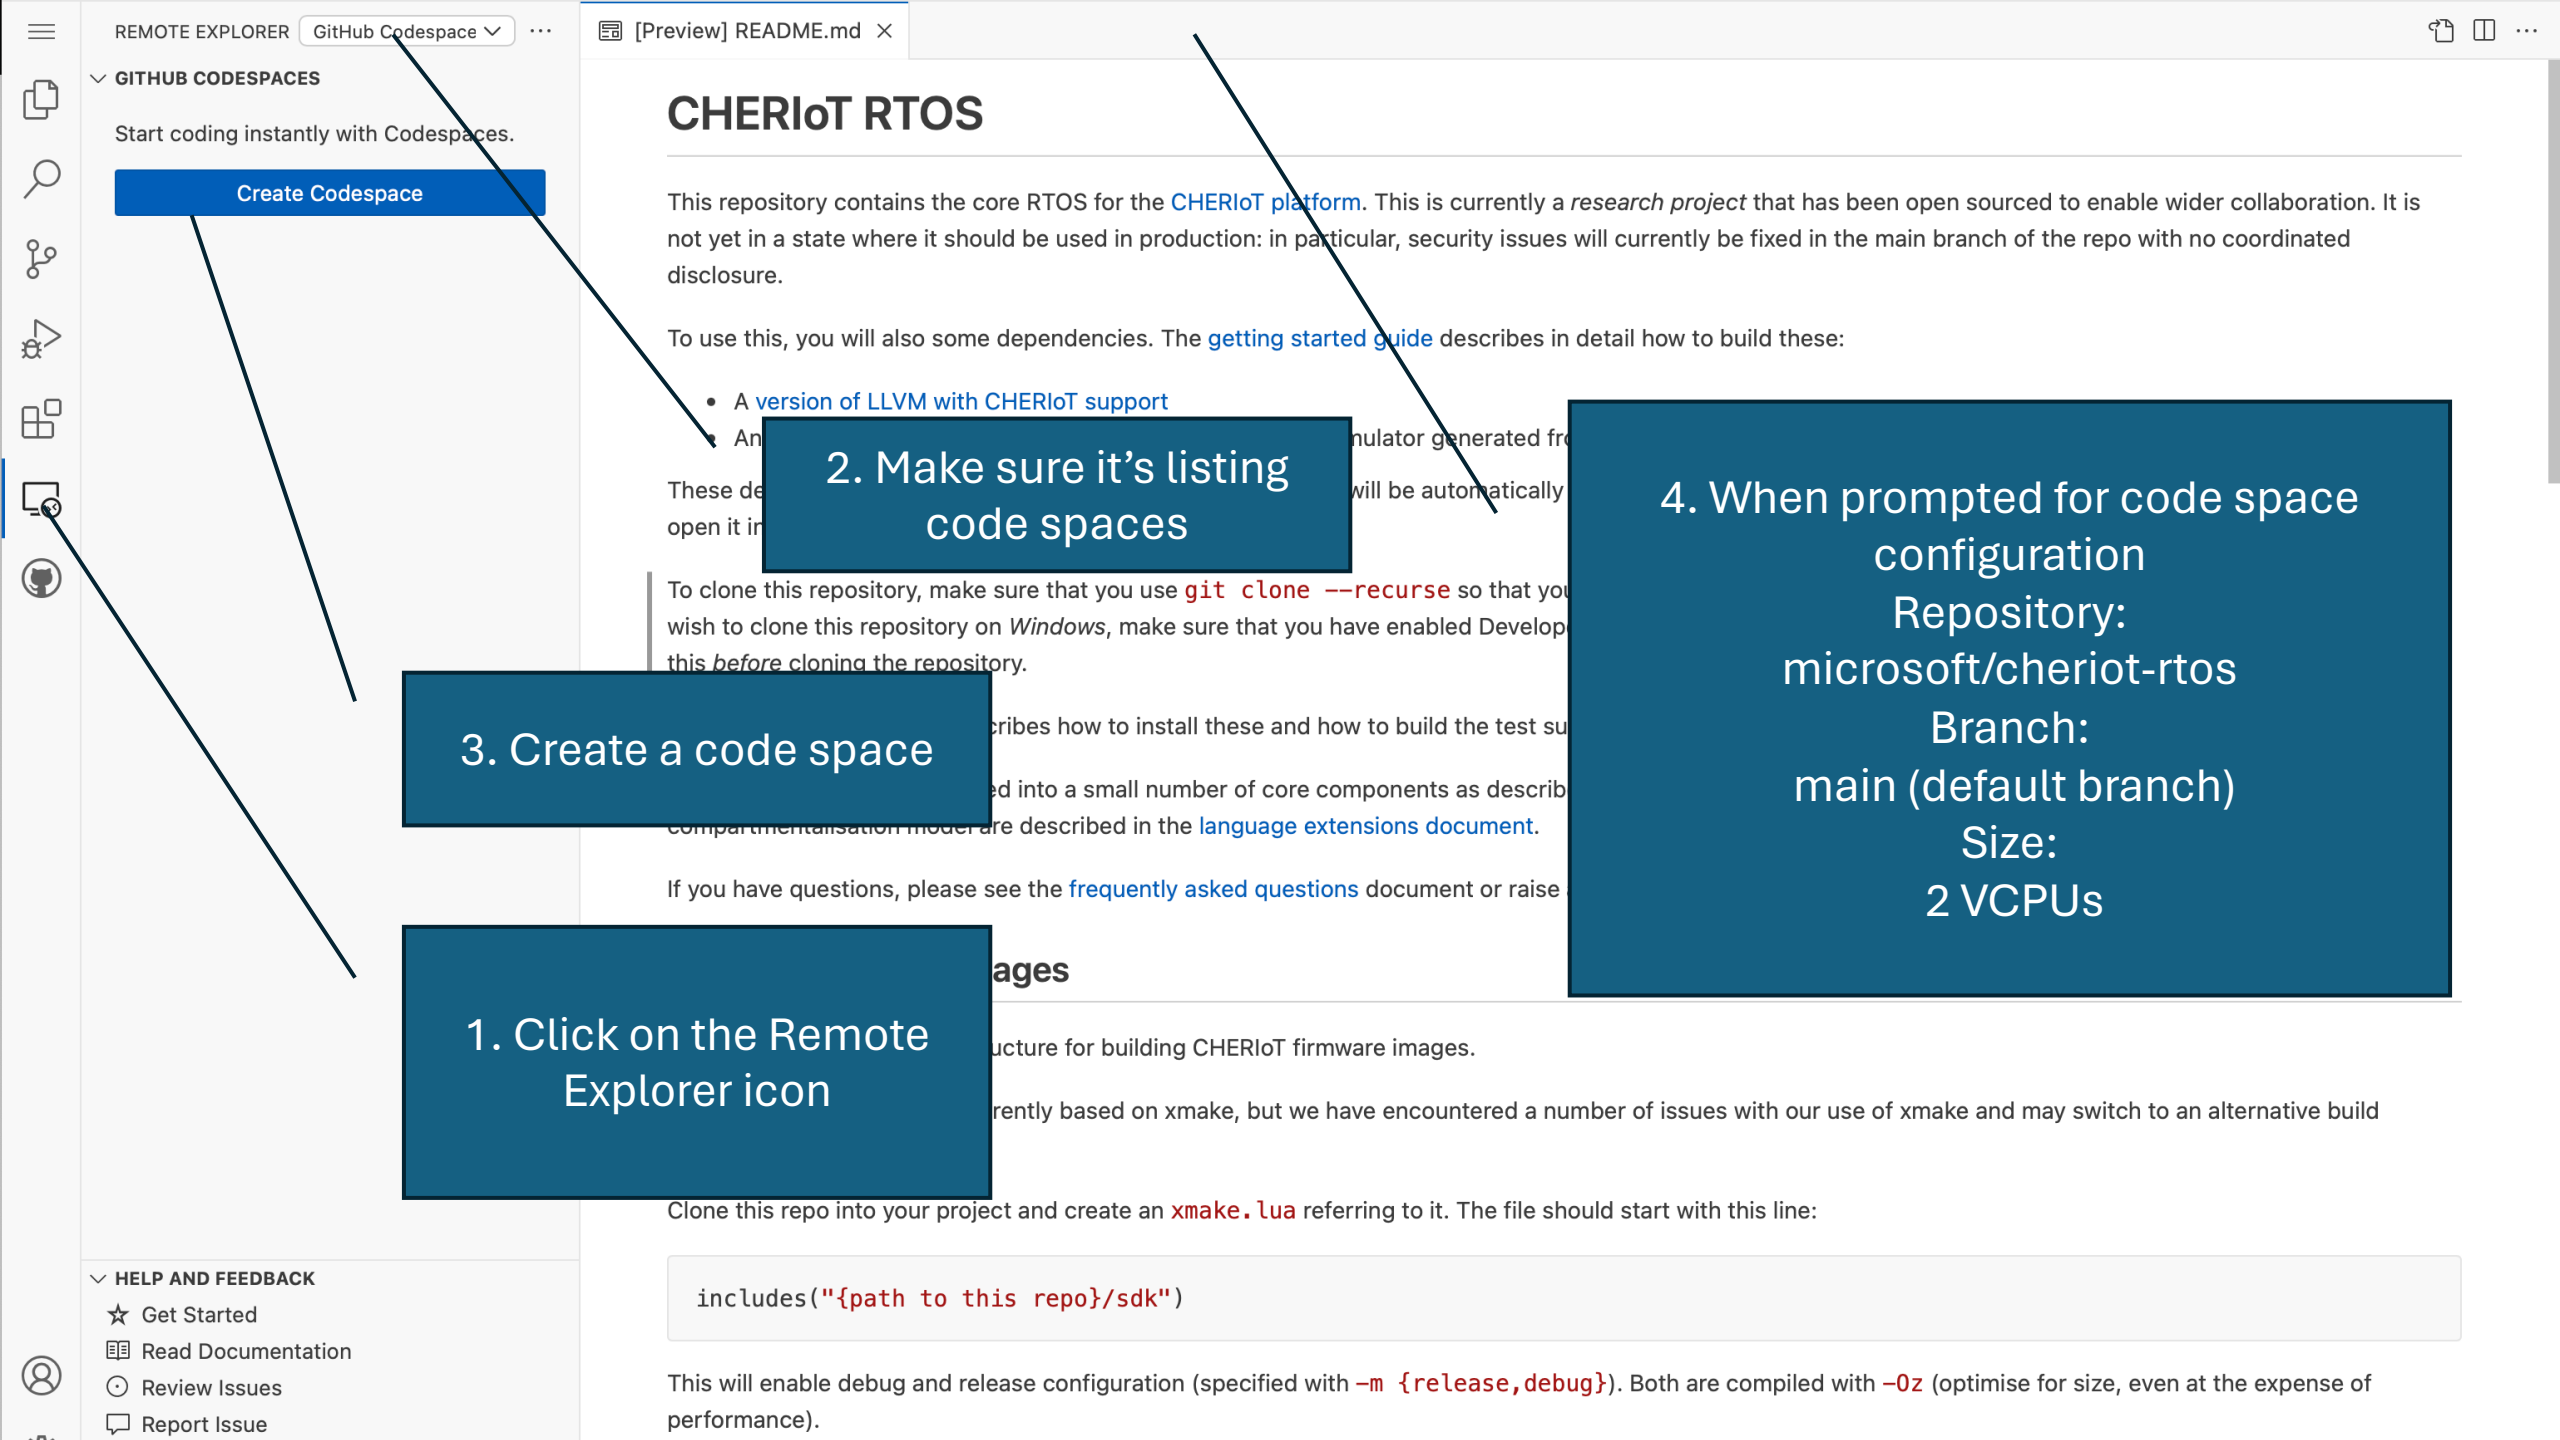
Add compartmentalization to C/C++

```
// Declaration adds an attribute to indicate
// the compartment containing the implementation
void __cheri_compartment("kv_store_sdk")
publish(char *key, uint8_t *buffer, size_t size);
```

```
-- Make sure it's compiled in the right
-- compartment in xmake.lua
compartment("kv_store_sdk")
    add_files("publish.cc")
    ...
```

Starting the exercise

<https://github.dev/microsoft/cheriot-rtos>



GITHUB CODESPACES

Start coding instantly with Codespaces.

Create Codespace

CHERIOT RTOS

This repository contains the core RTOS for the [CHERIOT platform](#). This is currently a *research project* that has been open sourced to enable wider collaboration. It is not yet in a state where it should be used in production: in particular, security issues will currently be fixed in the main branch of the repo with no coordinated disclosure.

To use this, you will also need some dependencies. The [getting started guide](#) describes in detail how to build these:

- A [version of LLVM with CHERIOT support](#)

These dependencies will be automatically installed when you open it in a Codespace.

To clone this repository, make sure that you use `git clone --recurse` so that you wish to clone this repository on *Windows*, make sure that you have enabled Developer Mode on your PC *before* cloning the repository.

This document describes how to install these and how to build the test suite. The repository is divided into a small number of core components as described in the [compartmentalisation model](#) are described in the [language extensions document](#).

If you have questions, please see the [frequently asked questions](#) document or raise an issue.

For more information on building CHERIOT firmware images, see the [building images](#) document.

Building images

The current build system is currently based on xmake, but we have encountered a number of issues with our use of xmake and may switch to an alternative build system.

Clone this repo into your project and create an `xmake.lua` referring to it. The file should start with this line:

Clone this repo into your project and create an `xmake.lua` referring to it. The file should start with this line:

```
includes("{path to this repo}/sdk")
```

This will enable debug and release configuration (specified with `-m {release,debug}`). Both are compiled with `-Oz` (optimise for size, even at the expense of performance).

HELP AND FEEDBACK

- ★ Get Started
- 📖 Read Documentation
- 🗨️ Review Issues
- 🗨️ Report Issue

3. Create a code space

1. Click on the Remote Explorer icon

2. Make sure it's listing code spaces

4. When prompted for code space configuration
Repository: microsoft/cheriot-rtos
Branch: main (default branch)
Size: 2 VCPUs

EXPLORER

- ✓ CHERIOT-RTOS [CODESPACES...]
- > .devcontainer
- > .github
- > benchmarks
- > docs
- > examples
- ✓ exercises / 01.compartmentalisation
 - JS cheri.js
 - JS crash.js
 - JS hello.js
 - js.cc
 - JS leak.js
 - \$ load_js.sh
 - microvium-ffi.hh
 - README.md
 - \$ run_simulator.sh
 - secret.cc
 - secret.h
 - xmake.lua
 - > scripts
 - > sdk
 - > tests
 - .clang-format
 - .clang-tidy
 - .gitignore
 - .gitmodules
 - ! azure-pipelines.yml
 - cgmanifest.json
 - CODE_OF_CONDUCT.md
 - compile_commands.json
 - compile_flags.txt
 - LICENSE
 - README.md

[Preview] README.md

CHERIOT RTOS

This repository is currently a *research project* that has been open sourced to enable wider collaboration. In particular, security issues will currently be fixed in the main branch of the repository.

Navigate to the compartmentalisation exercise.

To use the dependencies described in detail how to build these:

- A [version of LLVM with CHERIOT support](#)
- An implementation of the ISA (e.g. [CHERIOT-ibex](#) or the emulator generated from [the formal model](#)))

These dependencies are pre-installed in the dev container that will be automatically downloaded if you open this repository in Visual Studio Code or by hitting `.` to open it in GitHub Code Spaces.

To clone this repository, make sure that you use `git clone --recurse` so that you get all submodules. **IMPORTANT:** If you wish to clone this repository on *Windows*, make sure that you use `core.symlinks true`. You must do this *before* cloning the repository.

Split the terminal for side-by-side terminals

The [getting started guide](#) describes how to install these and how to build the test suite and examples in this repository.

The RTOS is privilege separated into a small number of core components as described in the [architecture document](#). The C/C++ extensions used by the compartmentalisation model are described in the [language extensions document](#).

If you have any questions or need help, please open an issue or raise an issue.

Read the README!

PROBLEMS

cheriot@co

bash + - [Terminal icons]

Split Terminal (⌘)

Run commands in the two terminals

Run the simulator in one terminal

```
$ cd exercises/01.compartmentalisation
$ xmake f --sdk=/cheriot-tools/
...
$ ./run_simulator.sh
...
JavaScript compartment: Secret stored at
0x2004cc8c (v:1 0x2004cc8c-0x2004cc90
l:0x4 o:0x0 p: G RWcgm- -- ---)
JavaScript compartment: Read 0x1ac bytes
of bytecode
JavaScript compartment: 0xdf8 bytes of
heap available
Hello world
```

Compile JavaScript and send it to the UART from the other

```
$ cd exercises/01.compartmentalisation
$ ./load_js.sh hello.js
Loading JavaScript:
...
Output generated: /dev/null
428 bytes
```

The exercise structure

- JavaScript code simulates an attacker with arbitrary code execution.
- Attacks from JavaScript can:
 - Leak a secret
 - Crash the system
 - Exhaust compartment memory

Each exercise will improve compartmentalisation to prevent one attack.

Three exercises



Move the code that owns the secret into a compartment.



Move the JavaScript interpreter into a compartment.



Prevent crashes from leaking memory in the JavaScript compartment.