

## M3 – Microkernel for Minimalist Manycores

Nils Asmussen

MKC, 06/19/2014

# Outline

- 1 Introduction
- 2 Spatial Isolation
- 3 Abstractions
- 4 Summary

# Outline

## 1 Introduction

## 2 Spatial Isolation

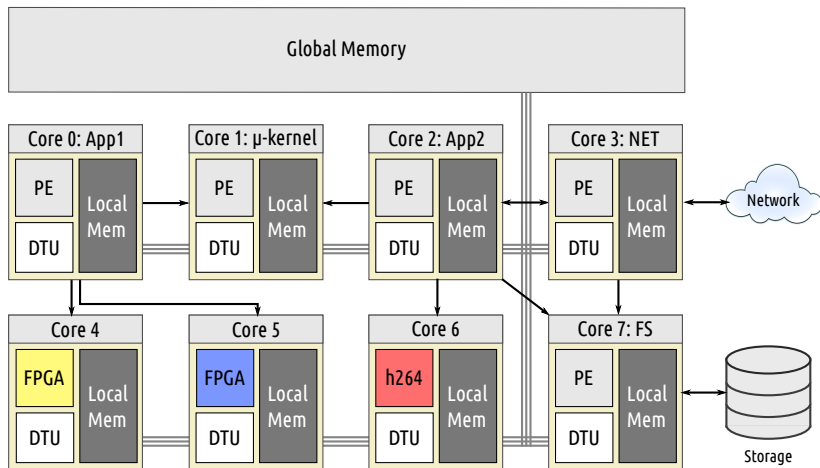
## 3 Abstractions

## 4 Summary

# Motivation

- We are moving towards accelerator architectures
- But how to really integrate heterogeneous cores and provide isolation?
- Energy consumption becomes more and more important
- Experiments show, that caches and MMUs consume lots of energy
- We try to address these issues by supporting heterogeneous cores, provide isolation, but be less energy demanding

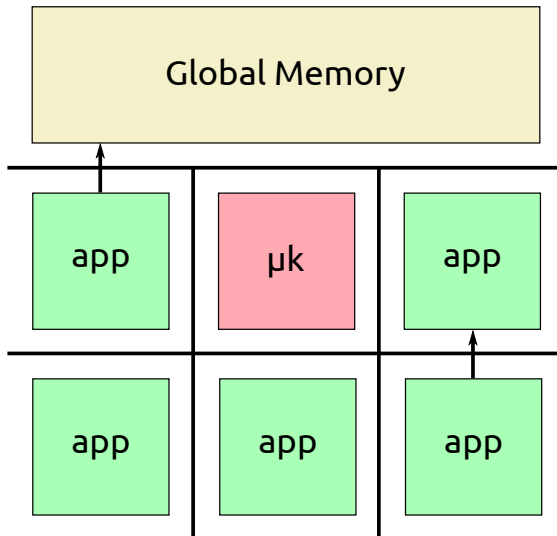
# Vision



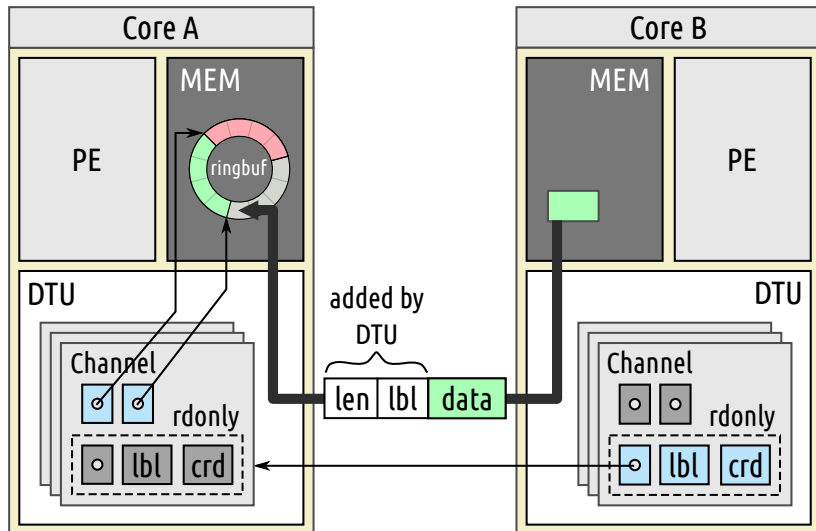
# Outline

- 1 Introduction
- 2 Spatial Isolation
- 3 Abstractions
- 4 Summary

# Isolation Principle

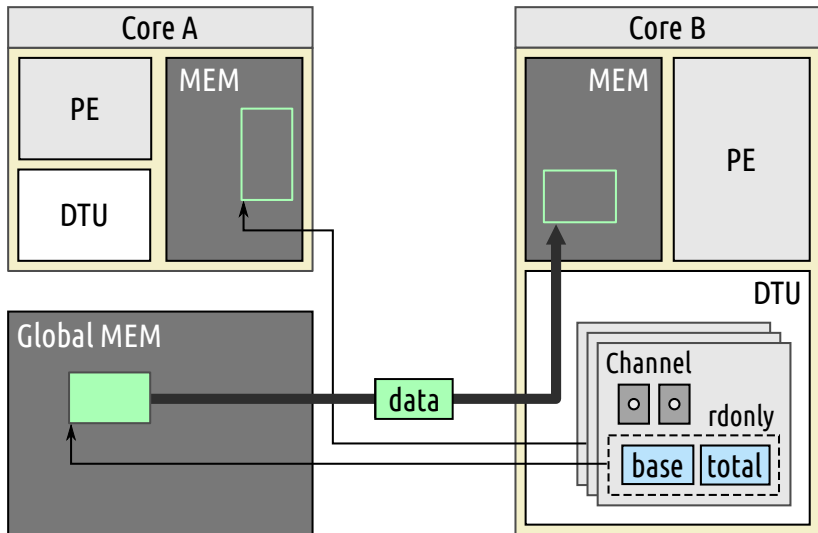


# Data Transfer between Cores





# Memory Access



# Outline

1 Introduction

2 Spatial Isolation

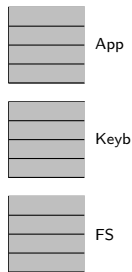
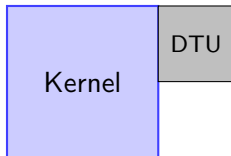
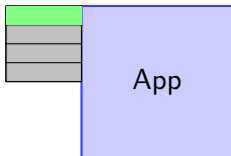
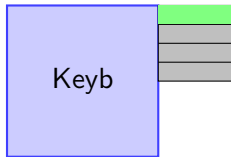
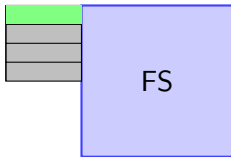
3 Abstractions

4 Summary

# Capabilities

- The  $\mu$ -kernel manages the capabilities of all tasks
- Like in L4, capabilities are protected by the kernel, selectors are managed by the applications
- Normal IPC can't be used because the kernel is not involved
- Thus, a special protocol is used to exchange capabilities

# Sessions



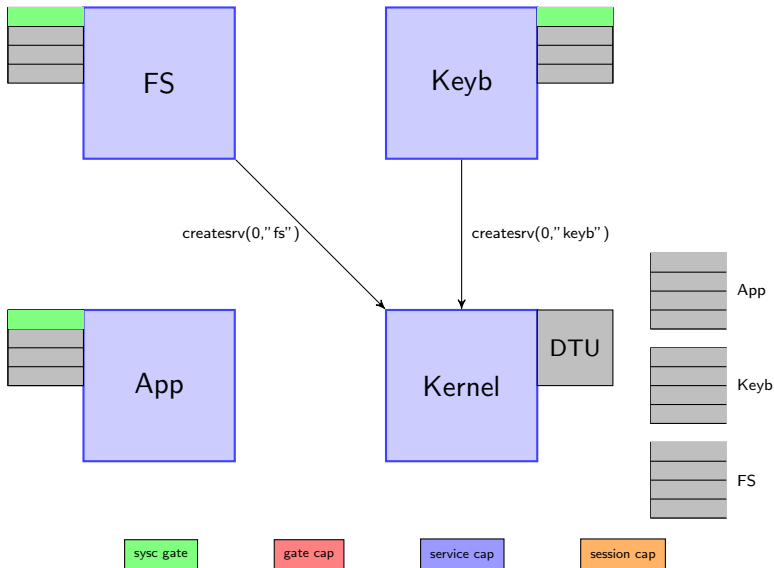
sysc gate

gate cap

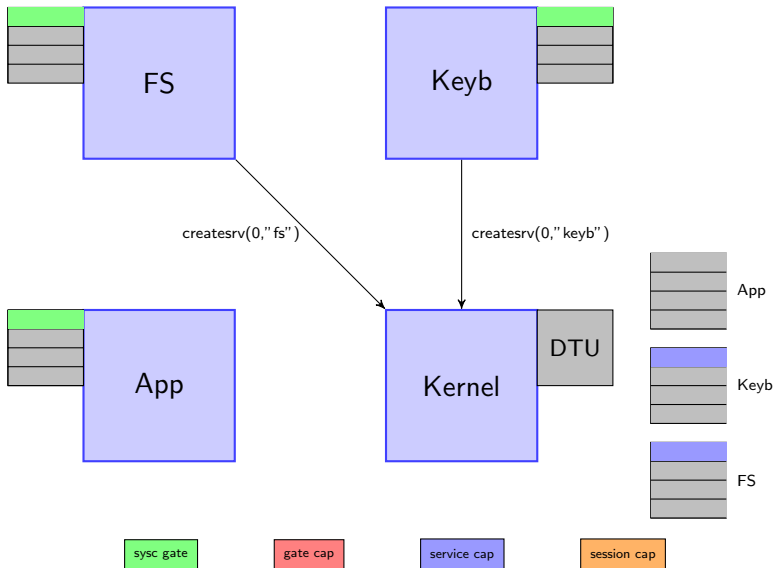
service cap

session cap

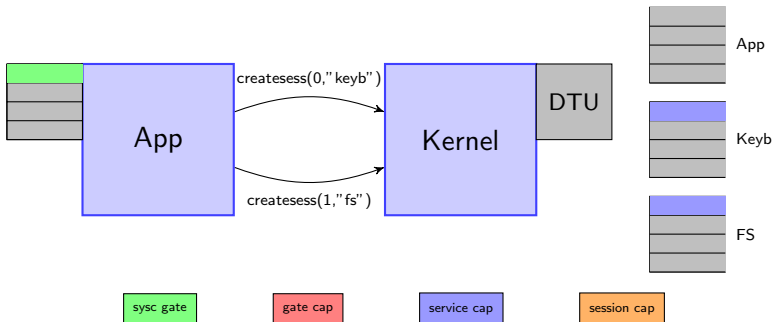
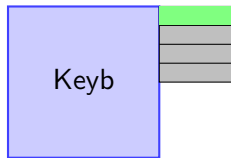
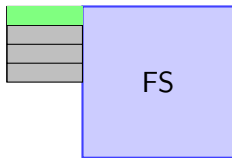
# Sessions



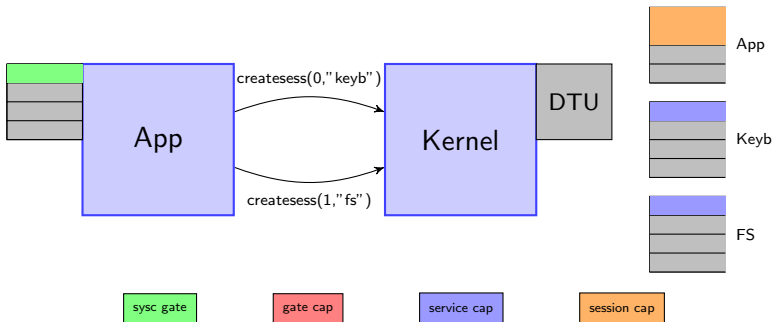
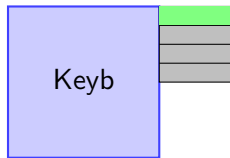
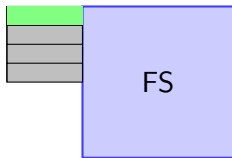
# Sessions



# Sessions

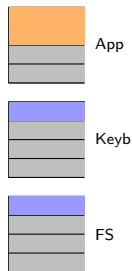
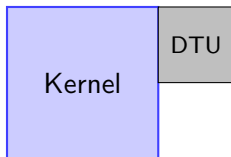
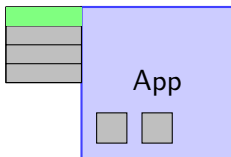
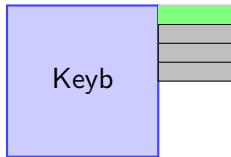
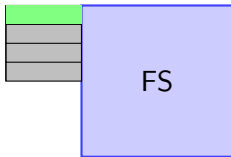


# Sessions





# Sessions



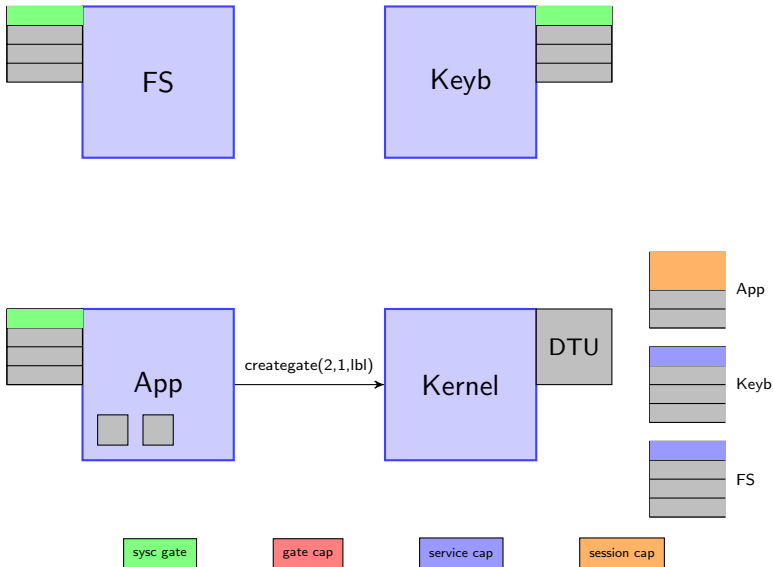
sysc gate

gate cap

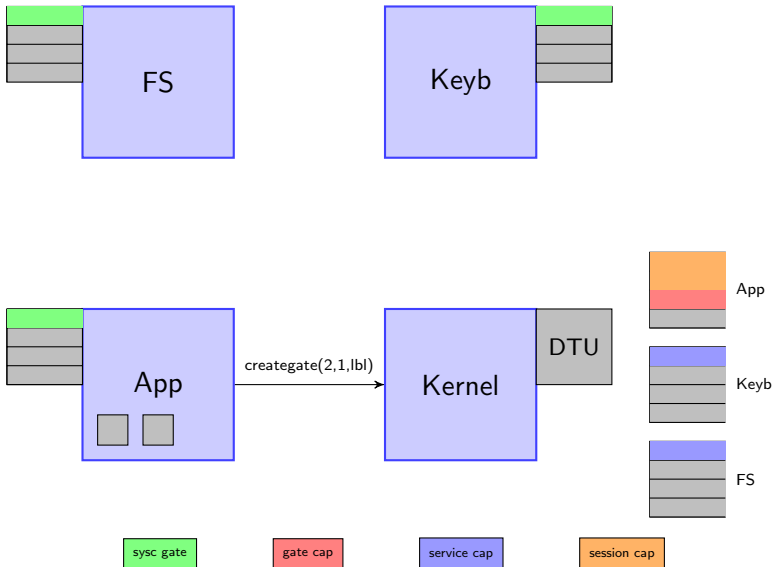
service cap

session cap

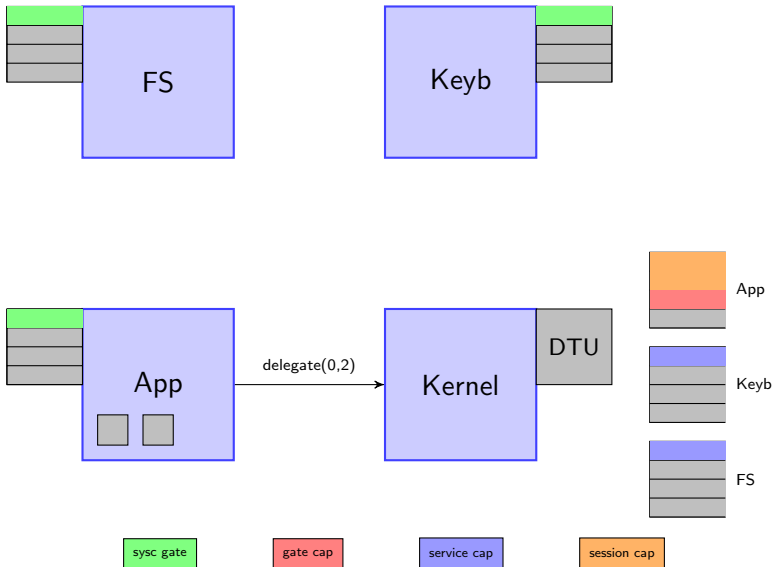
# Sessions



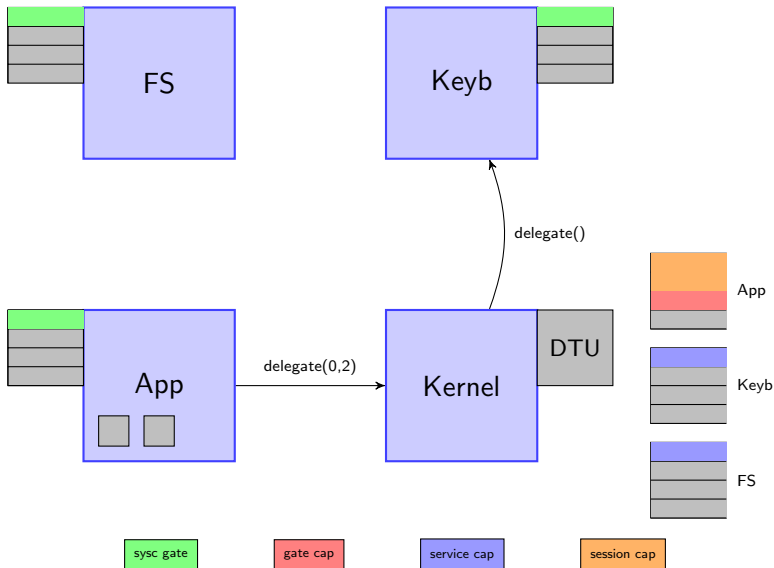
# Sessions



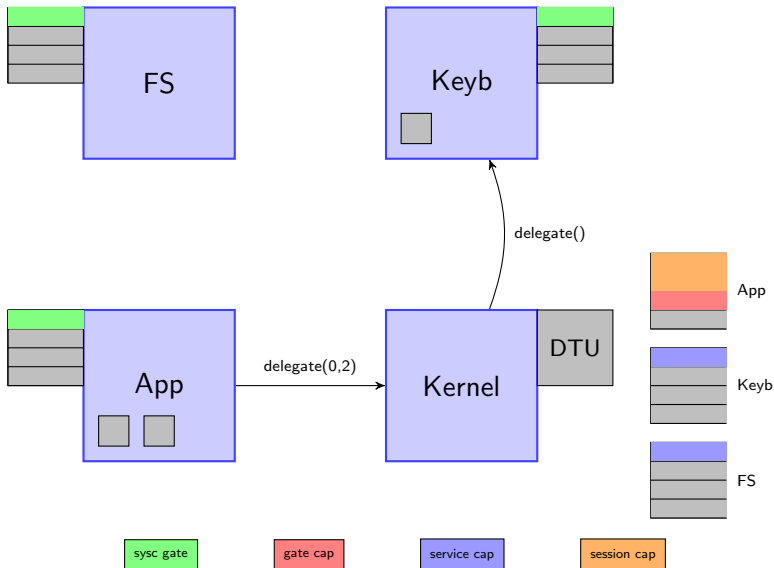
# Sessions



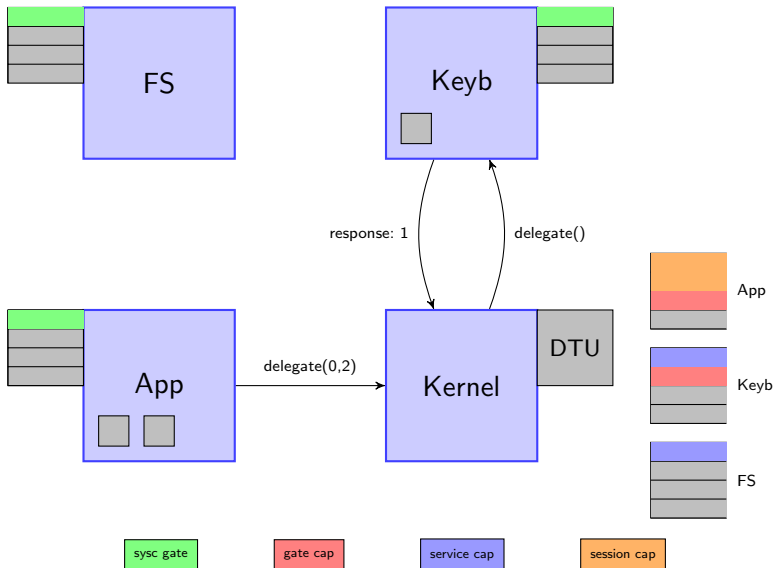
# Sessions



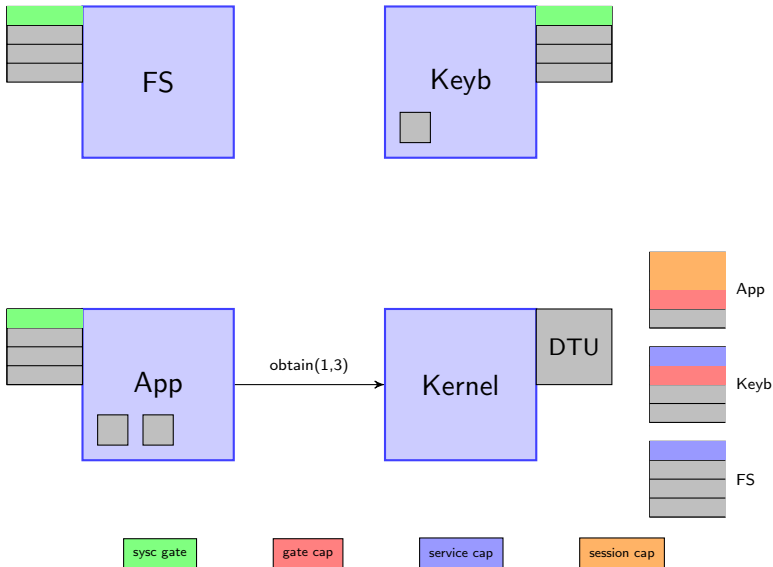
# Sessions



# Sessions

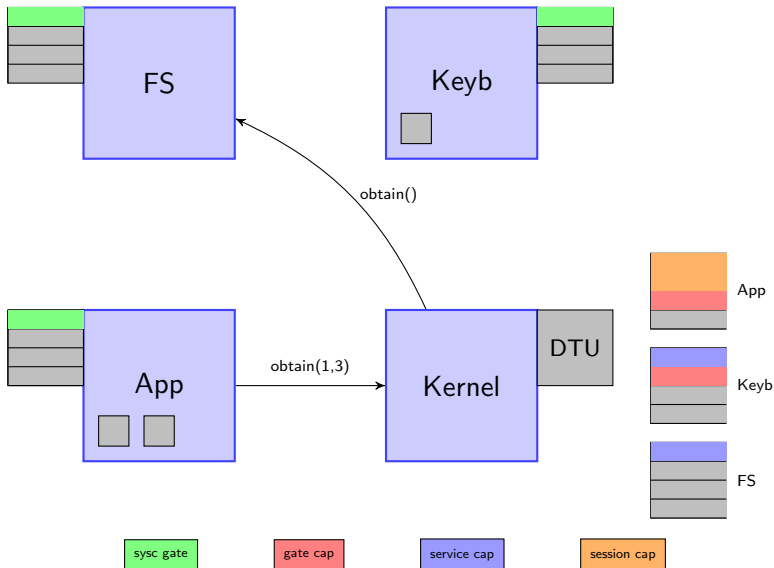


# Sessions

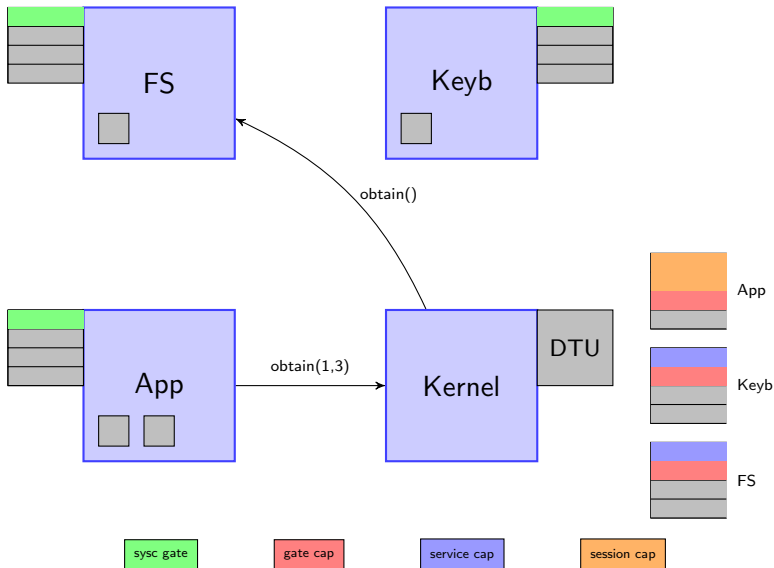




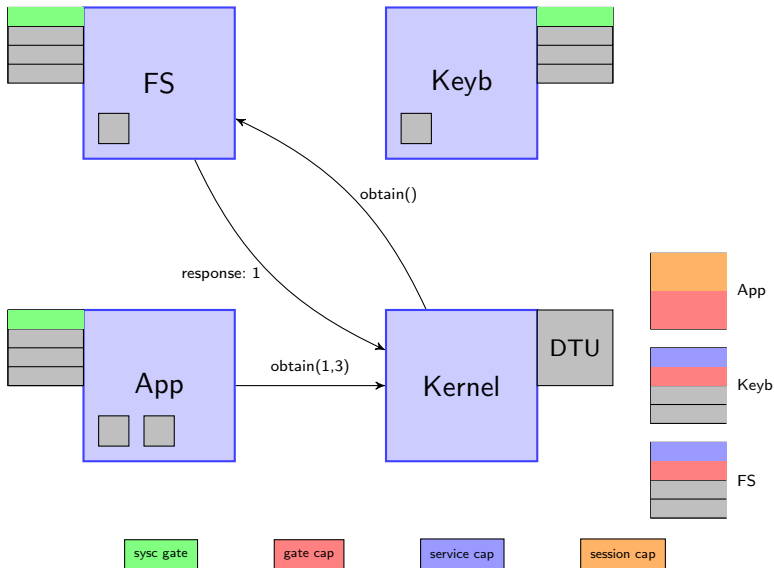
# Sessions



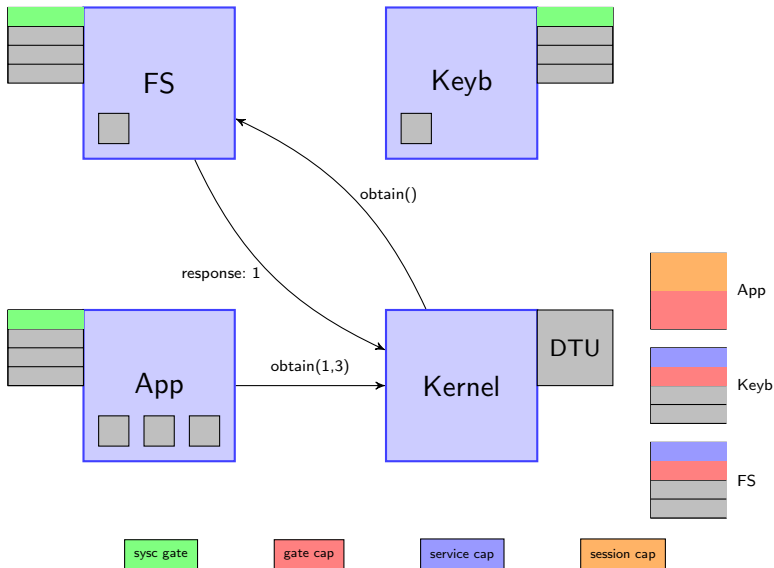
# Sessions



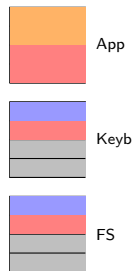
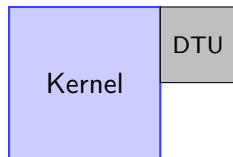
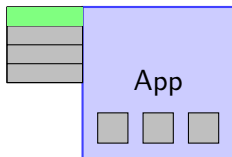
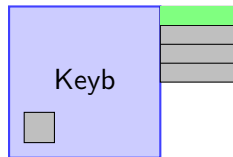
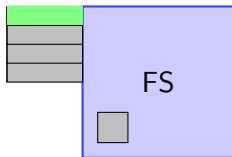
# Sessions



# Sessions



# Sessions



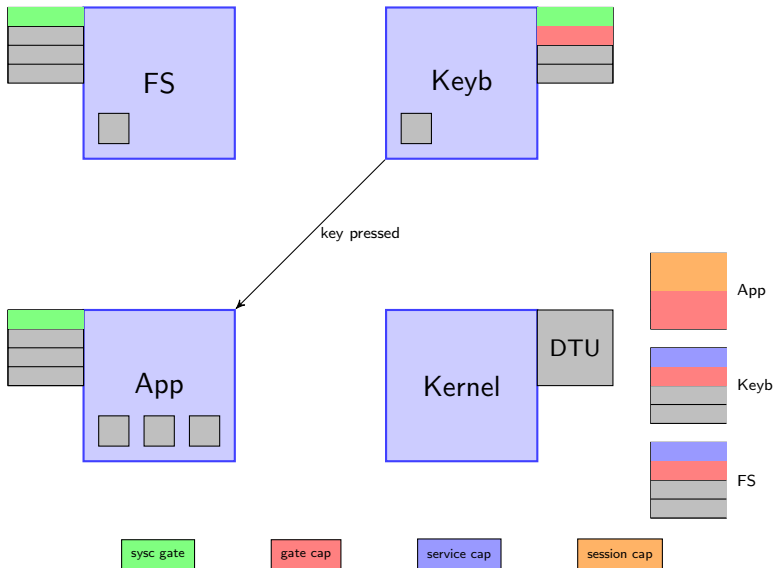
sysc gate

gate cap

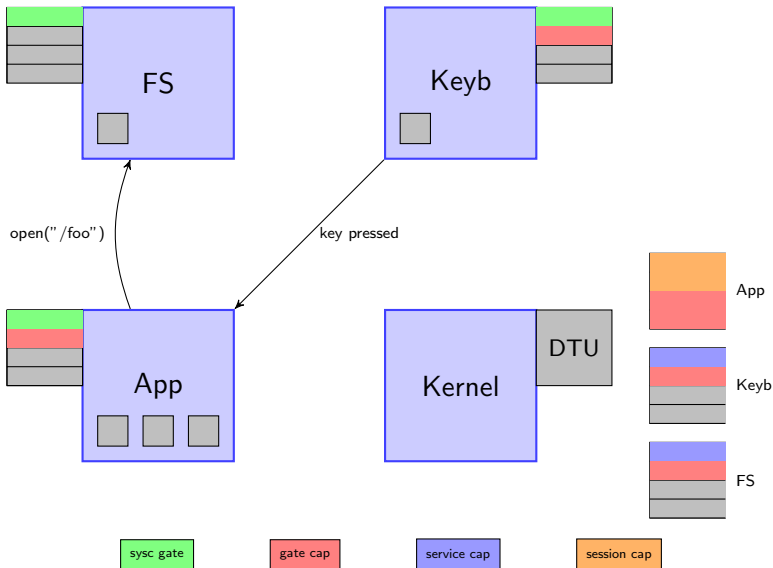
service cap

session cap

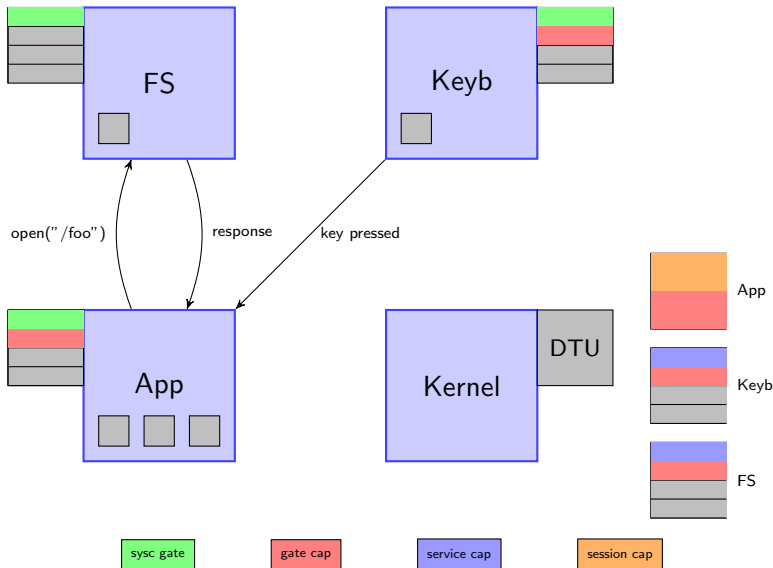
# Sessions



# Sessions



# Sessions





# Memory

- The syscall `reqmem` can be used to request global memory
- Global memory is managed by the kernel
- This memory capability can be put on a channel to access it
- For that reason, memory capabilities are the same as gate capabilities
- I.e. unlike L4, there are no different kinds of capabilities
- Syscall `derivemem` to create a memory capability for a subset of a given one
- Syscall `createmem` to create a memory capability for the own local memory

# Filesystem

- There is a filesystem service running on some core
- The FS uses global memory as the inode and buffer cache
- App goes to FS in order to know where file fragments are in global memory
- FS hands out memory capabilities to the application
- App reads/writes directly from/to global memory afterwards

# Filesystem – API

- Like in L4Re, the mount-tree is application local
- The File abstraction is the low level abstraction to read/write files
- The BufferedFile abstraction sits on top and provides a buffer in between

# Filesystem – API

- Like in L4Re, the mount-tree is application local
- The File abstraction is the low level abstraction to read/write files
- The BufferedFile abstraction sits on top and provides a buffer in between

## Example

```
VFS::mount("/", "m3fs");

BufferedFile file;
VFS::open("/myfile", VFS::R, file);
while(file.read(buf, 512) > 0) {
    // ...
}
```

# Tasks

- Syscall `createtask` to create a new task
- The kernel assigns the task to a core, if there is a free one
- The application gets the task capability and a memory capability for the whole local memory of that core
- Thus, it has complete control over that core
- Syscall `start` to start the task
- It can exchange capabilities with the task before starting it
- Syscall `wait` to wait until a task is finished (called `exit`)

# Tasks – Examples

## Executing ELF-Binaries

```
Task task("test");  
char *args[] = {"/bin/hello", "foo", "bar"};  
task.exec(3, args);
```

# Tasks – Examples

## Executing ELF-Binaries

```
Task task("test");  
char *args[] = {"/bin/hello", "foo", "bar"};  
task.exec(3, args);
```

## Asynchronous Lambdas

```
Task task("test");  
int a = 10;  
task.delegate(CapRngDesc(1, 2));  
task.run([a]() {  
    Serial::get() << "a=" << a << "\n";  
});
```

# Outline

1 Introduction

2 Spatial Isolation

3 Abstractions

4 Summary



# Summary

- M3 is a different kind of microkernel-based system
- The  $\mu$ -kernel manages remote resources instead of local resources
- The  $\mu$ -kernel and applications run alone on the cores
- Applications are linked against a library to get abstractions
- Capabilities are used to manage permissions in the system
- Similar operations with caps, but still different