

Caps²: L4Re and CHERI

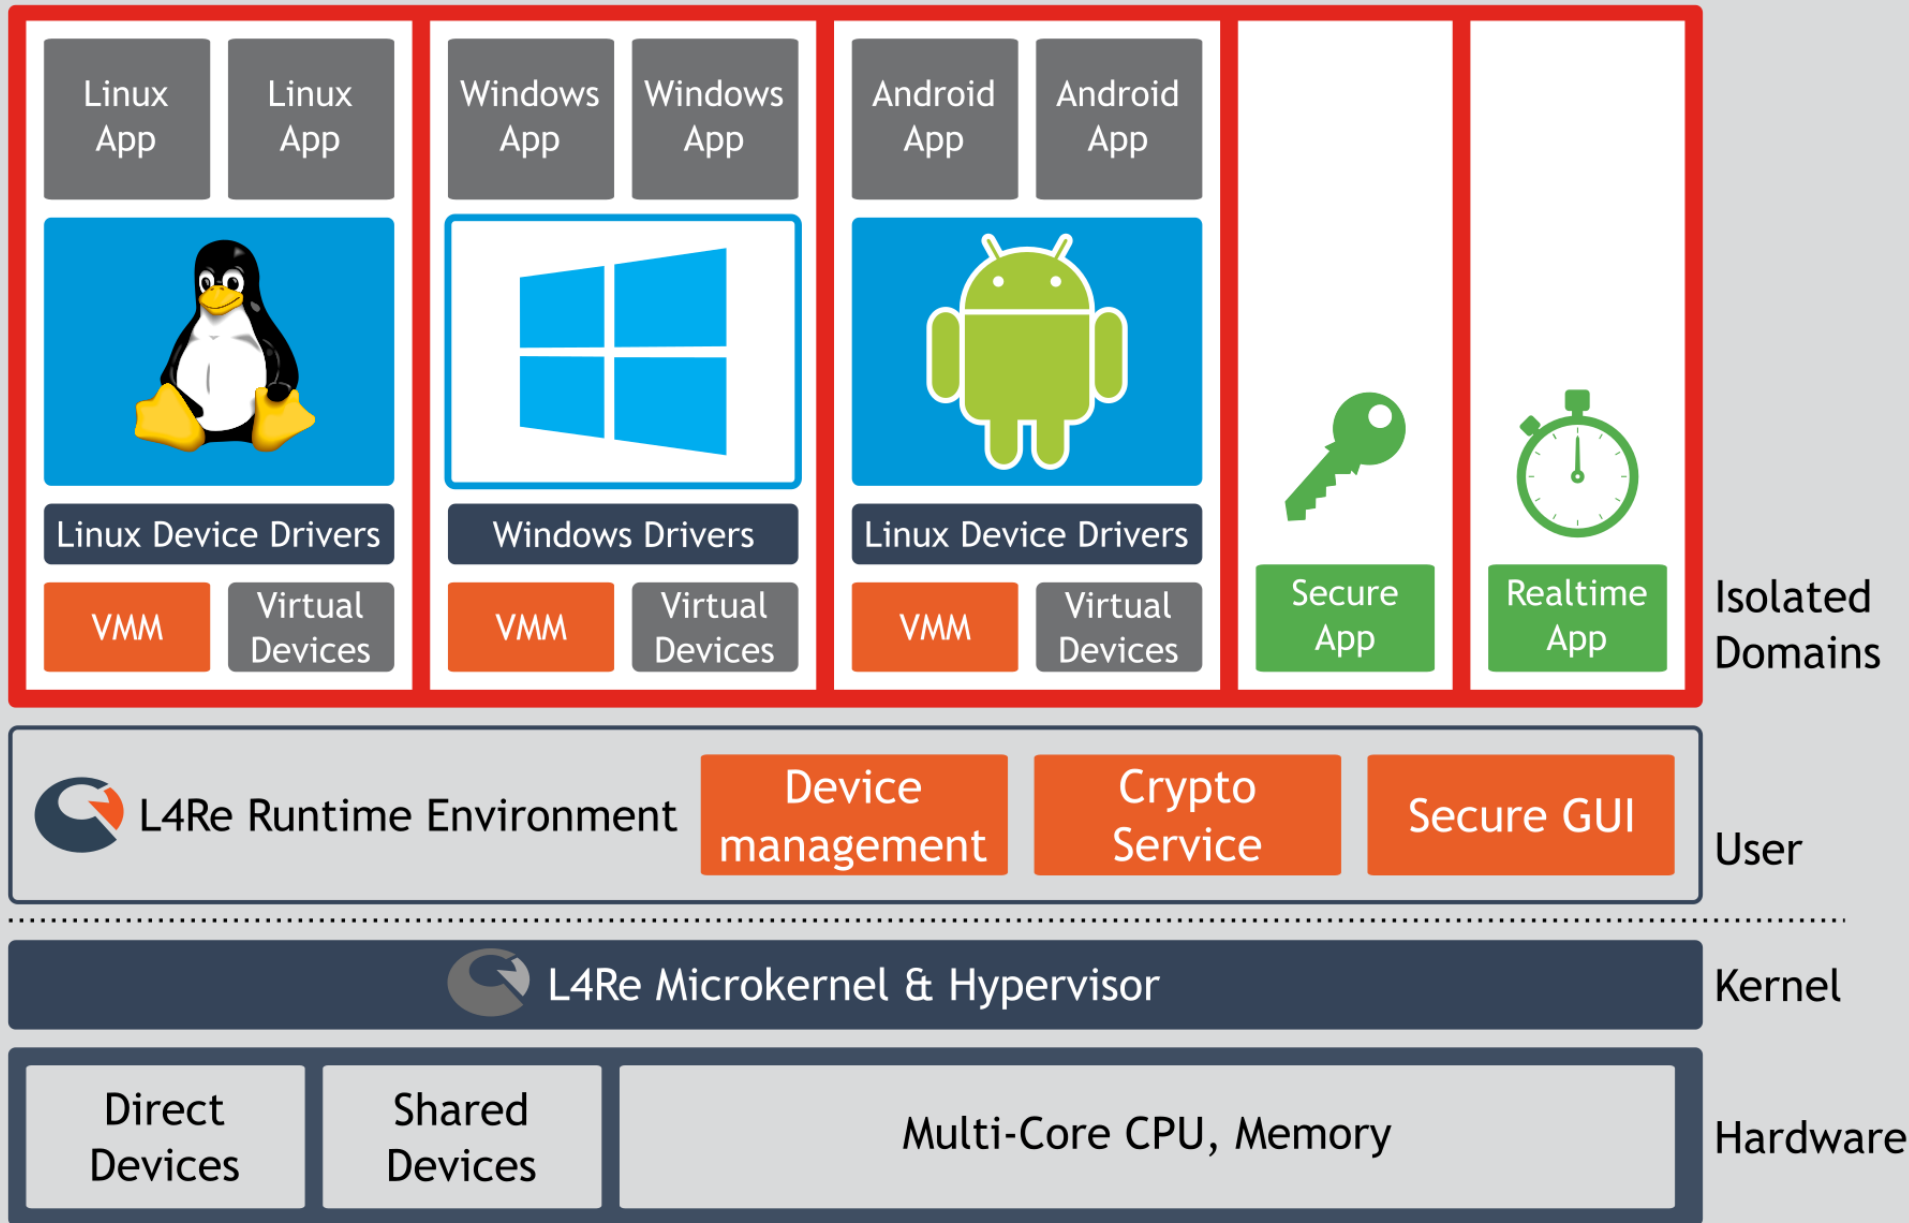
Adam Lackorzynski, Hermann Härtig

Apr 23rd 2016



- L4 microkernel system
- OS framework – Security, RT, Virtualization
- Developed since 1996
- Small TCB, component based, isolating
- Capability-based object system: Kernel & User
- SMP – ARM, x86, MIPS – VT, SVM, HYP, TZ, VZ
- Paravirtualization (e.g. L⁴Linux)
- Libc, C & C++, pthreads, POSIX, client/server framework, platform management
- Many (ARM) platforms
- Commercially supported & open source

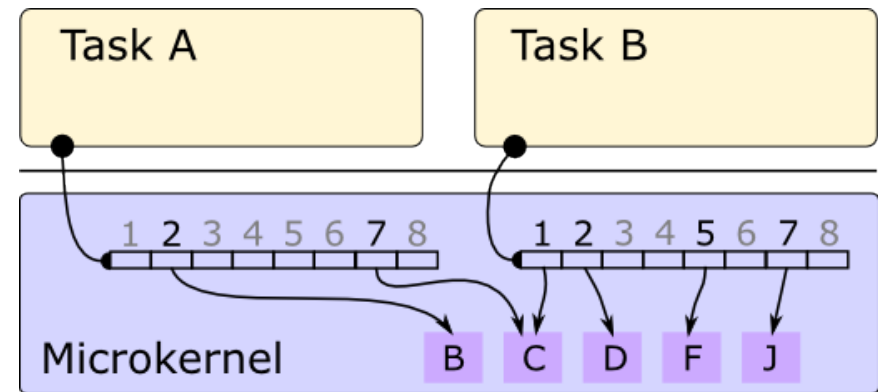
L4Re Architecture Overview

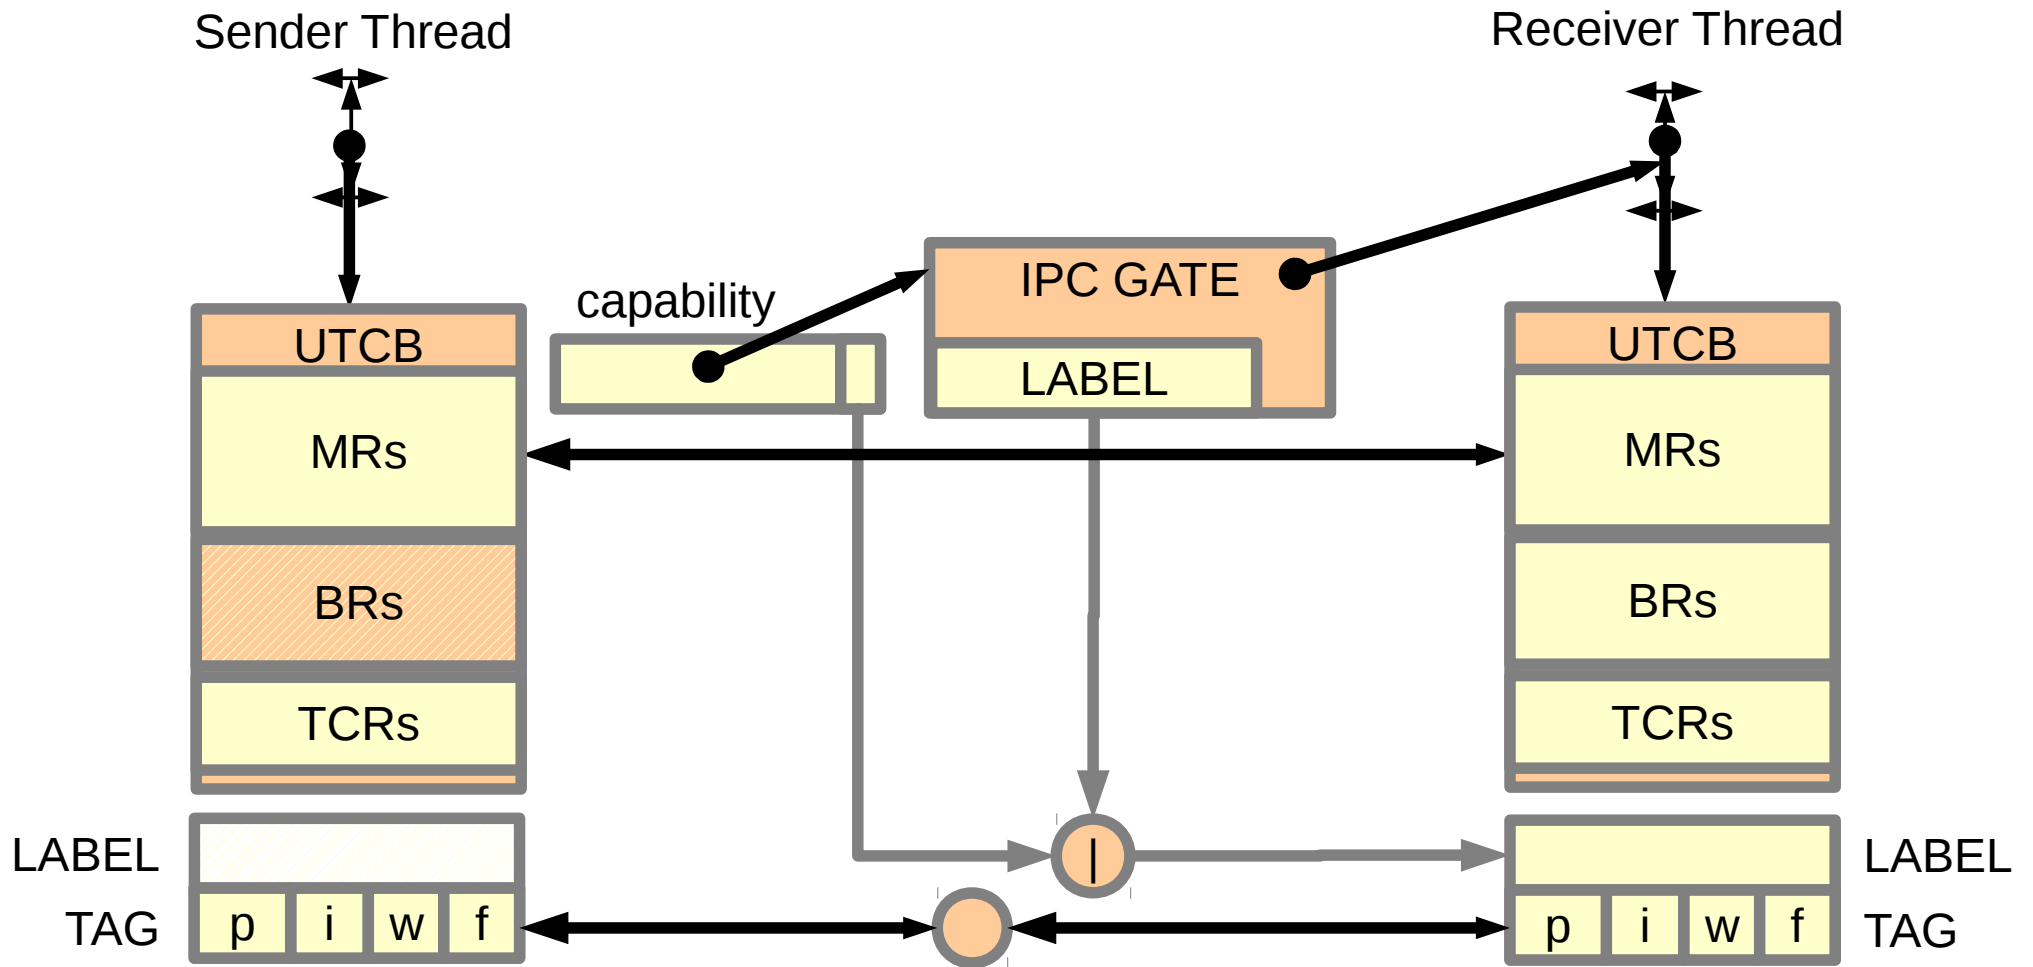


- Everything is an object
 - All system calls run on an object
 - „one system call“ - send message to object
- Reference to object: capability
- Uniform interfaces
 - For kernel and user-level objects

- Pass access rights to physical memory pages to other domains
 - Flex-page: Size-aligned region description
- Revoke access rights
- Recursively: Stored in kernel

- References to kernel objects
 - Local naming (per protection domain)
 - Access control
- Kept in per task capability table, managed by user-level
- Pointer sized
- Can be passed to other protection domains
- User-level objects:
 - Via Kernel-provided communication object
- IPC sends to == invokes a capability

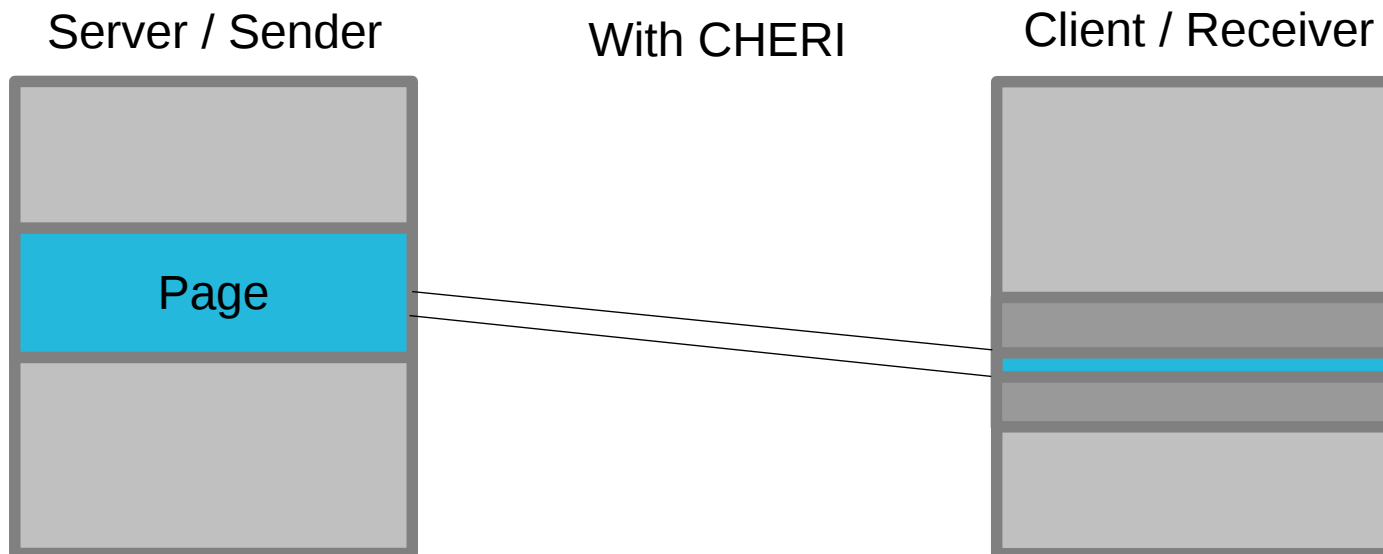
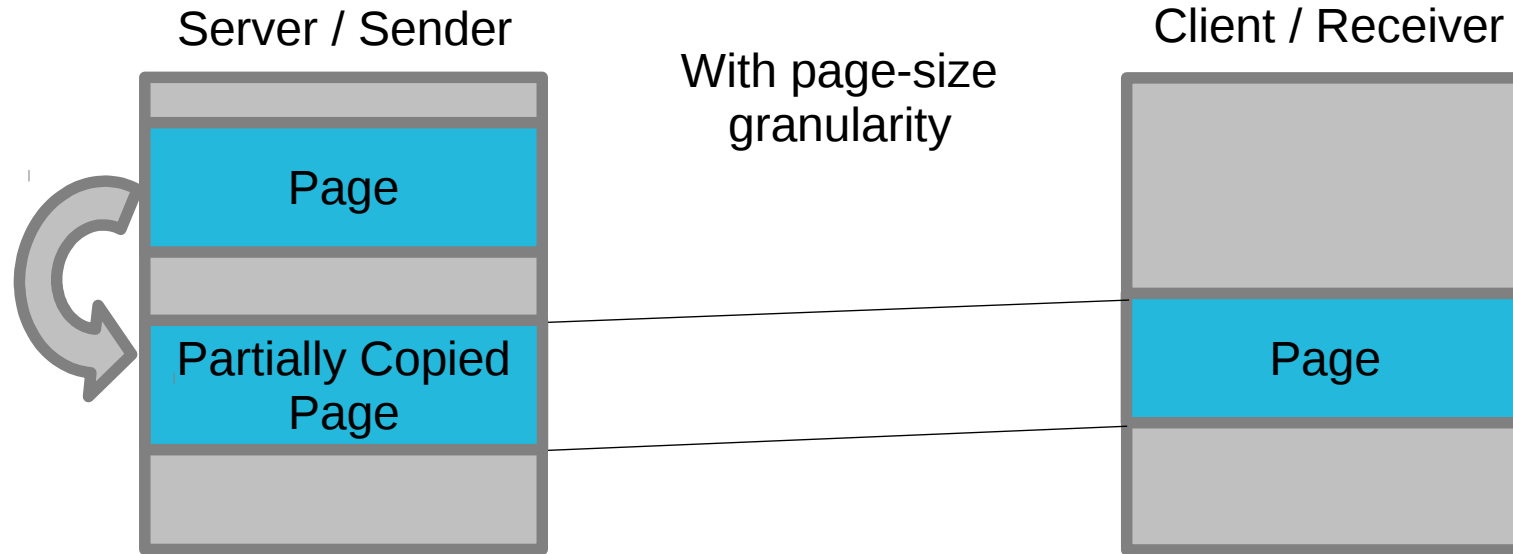




- Page faults & other exceptions
- Each thread has handler caps
 - Upon faults, the kernel invokes those caps
 - In good case: Handler replies with a resolving IPC containing a memory mapping / new CPU state
- Memory mappings can also be made explicit
- Memory access granularity: 4k (typically)

- Pass Cheri-Cap(s) via L4 messages to transfer fine-granular access rights
- Cheri-Cap faults as L4 exceptions

Byte-level Sharing



Fine-grained memory access:

- No more copy-in/copy-out from data structures
- Sharing densely packed data
- Fine-grained access to MMIO

Applications:

- Device drivers (but could (should!) be fixed by HW)
- Database: KVS, ...
- ...

l4re.org