


Beyond Address Spaces -

Flexibility, Performance, Protection, and Resource Management in the Type-Safe JX Operating System

Michael Golm

University of Erlangen-Nuremberg
Computer Science Department 4
(Operating Systems and Distributed
Systems)
Martensstraße 1
91058 Erlangen, Germany



Overview

- Objectives
- Architecture
- Protection
- Flexibility
- Performance

Objectives of the JX system

- Dynamic OS extension with untrusted components
- Code reuse
 - tailored OS configurations
- Exact resource accounting and customizable management
- Protection
 - flexibility
 - performance
 - robustness

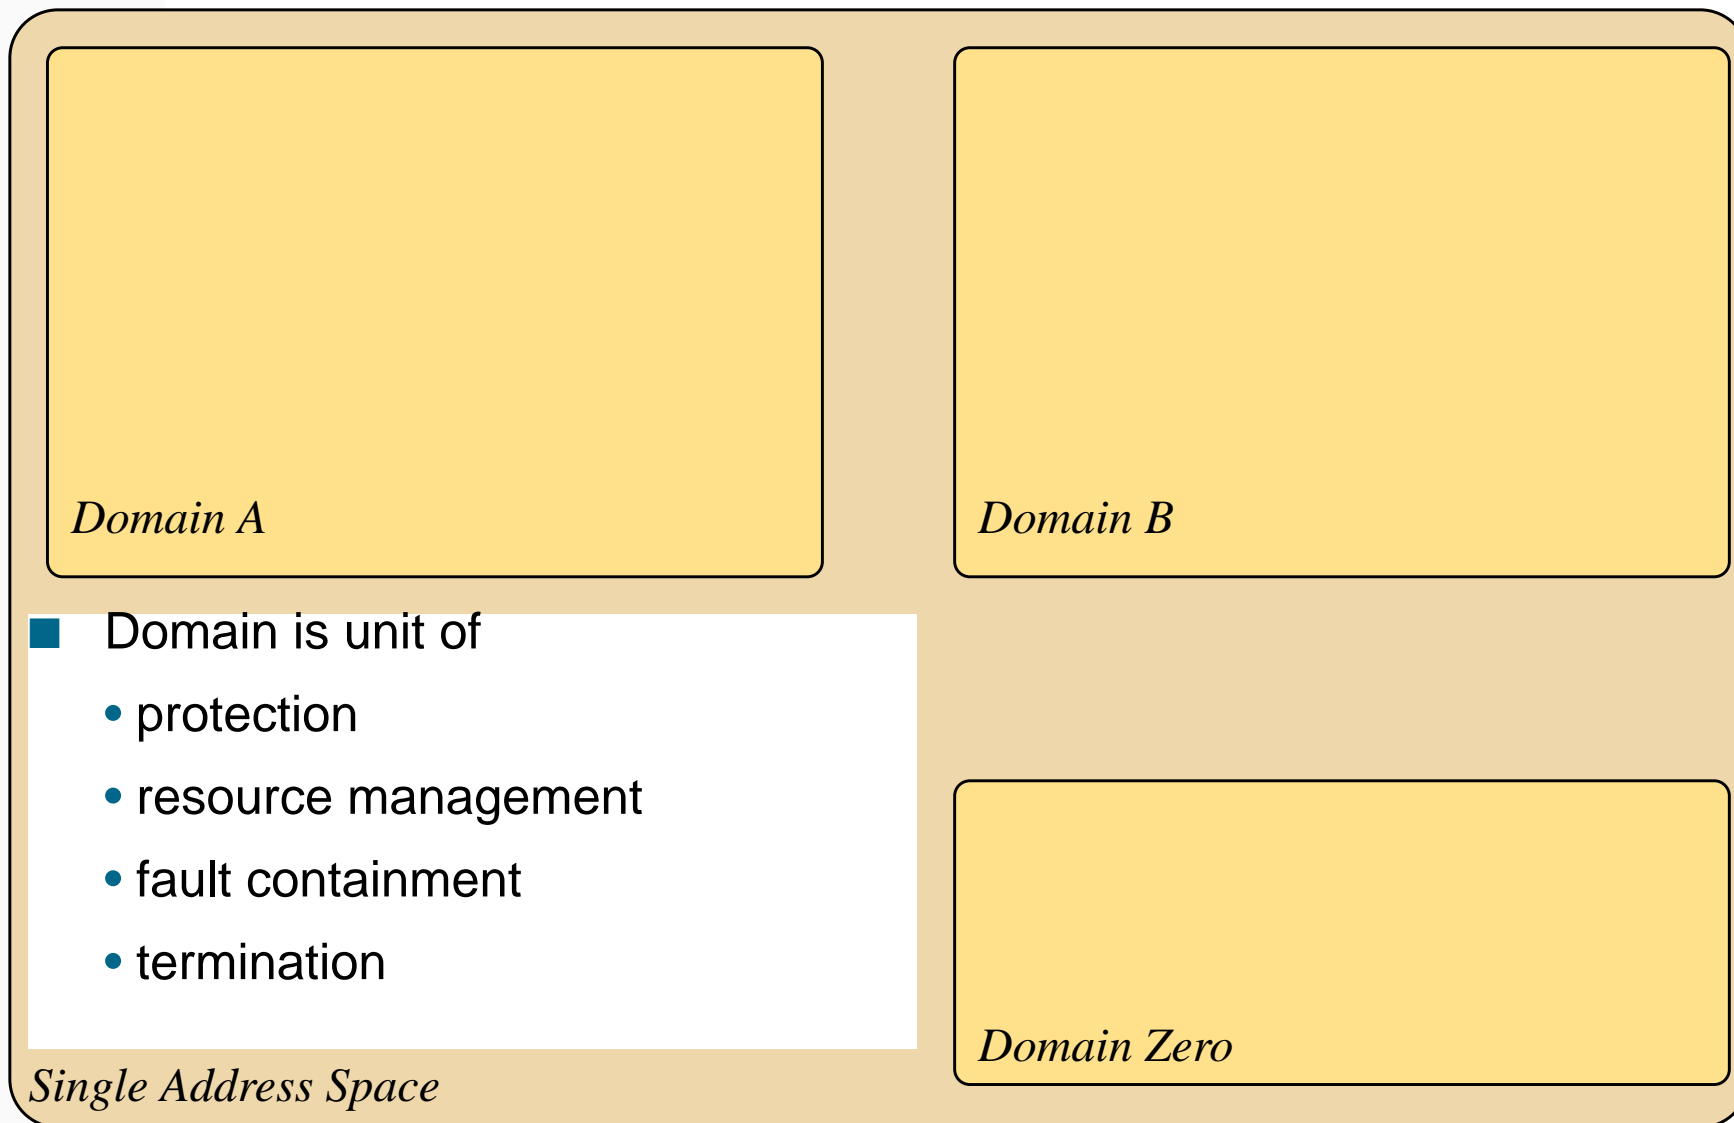
Single Address Space

JX Architecture

Single Address Space

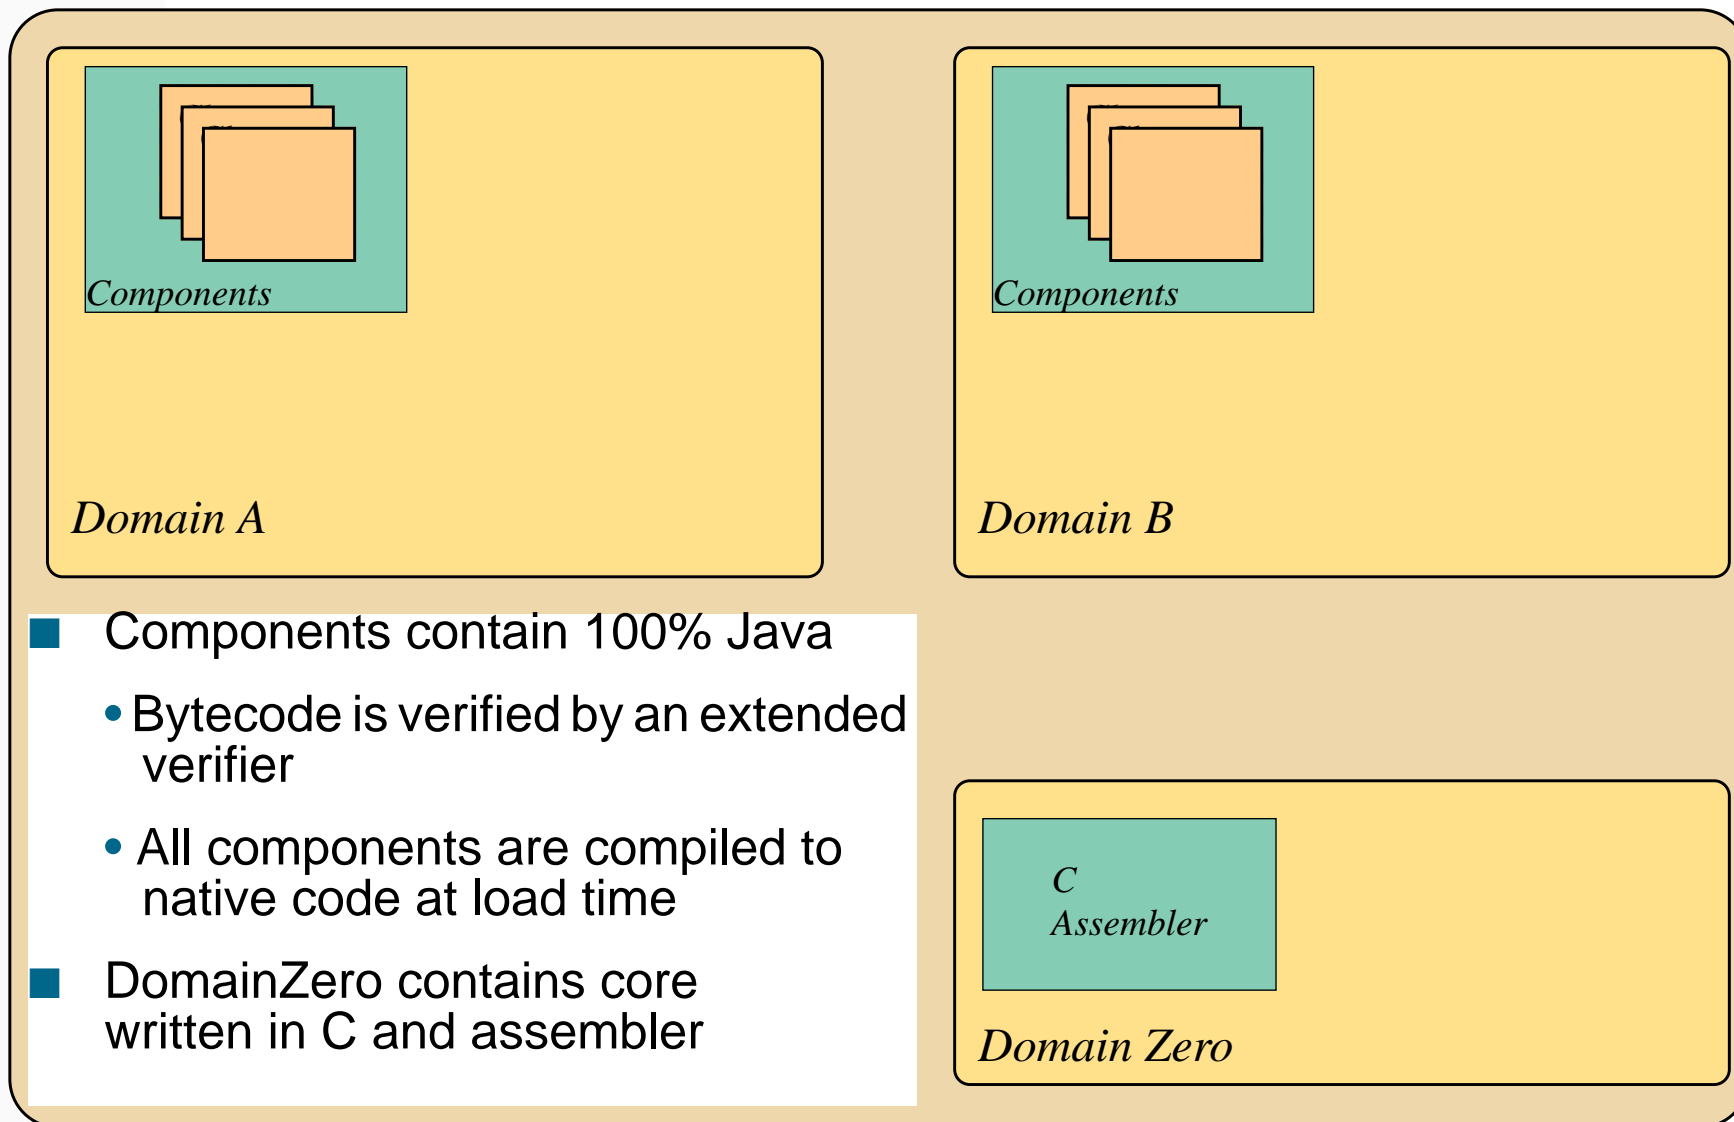
Domains

JX Architecture



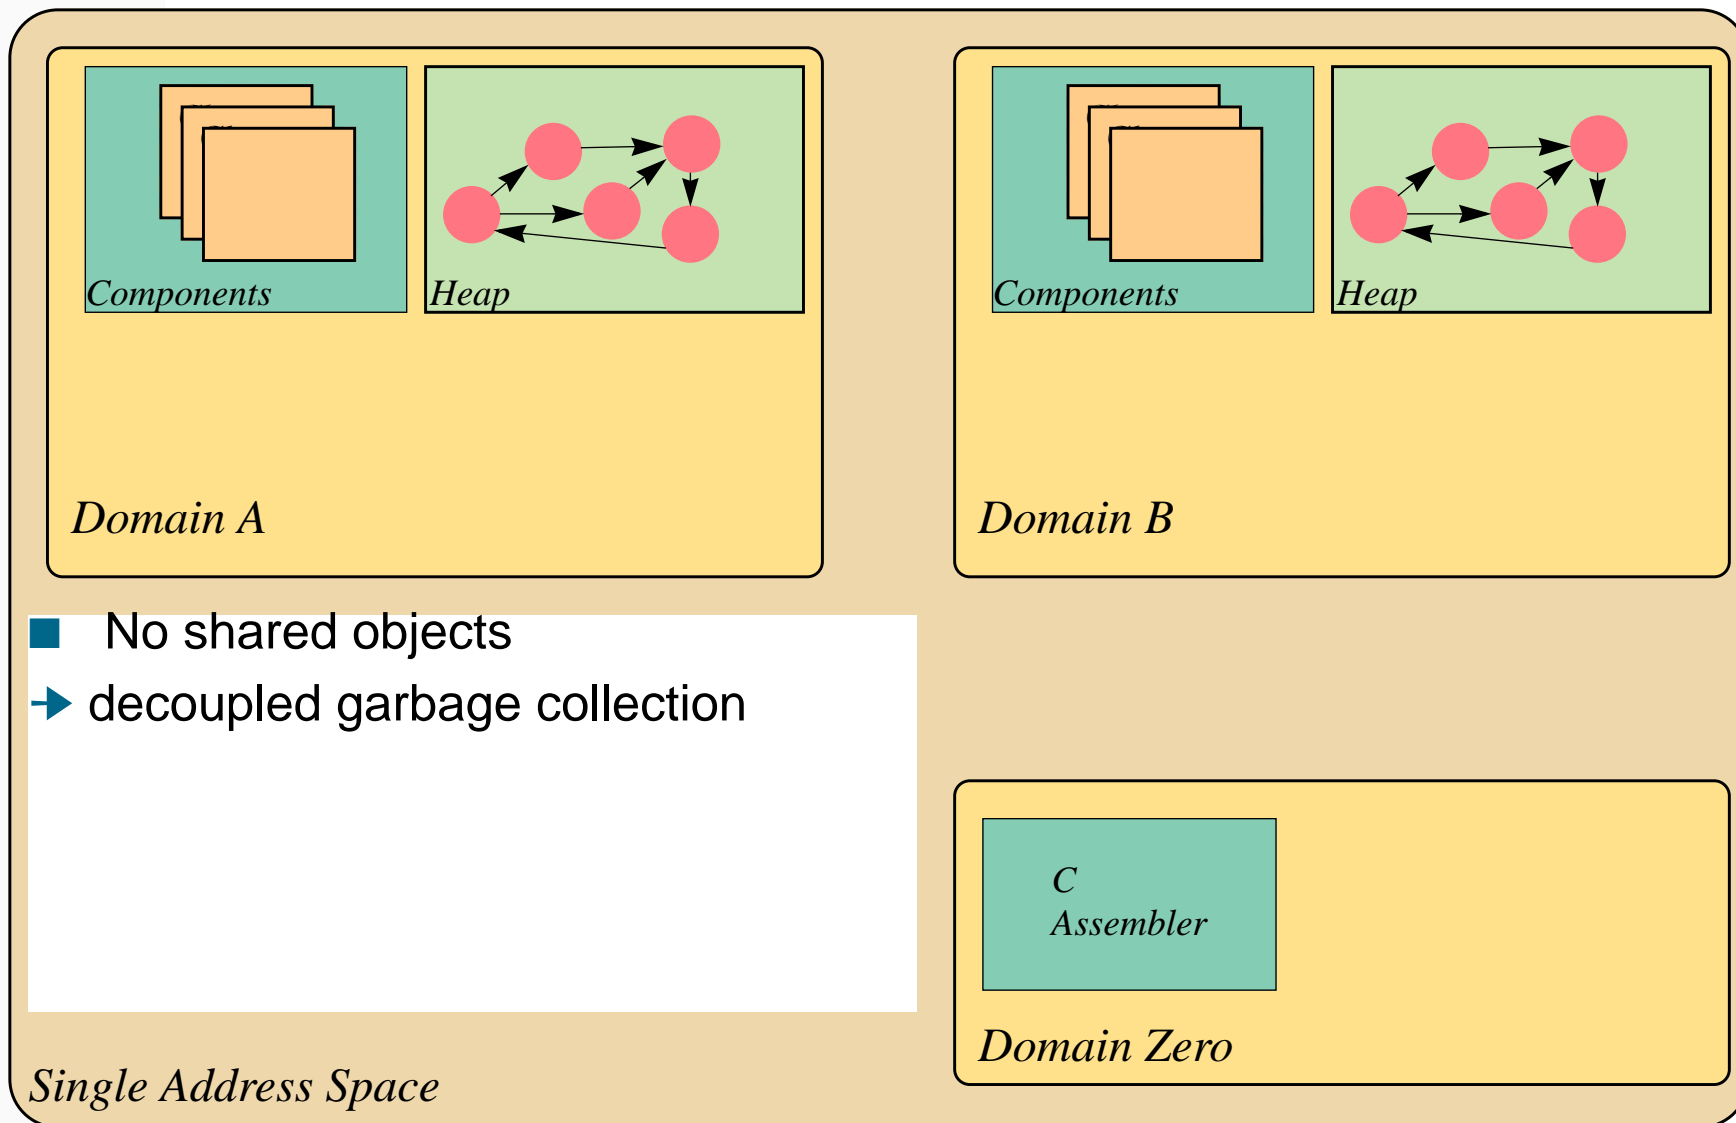
Components

JX Architecture



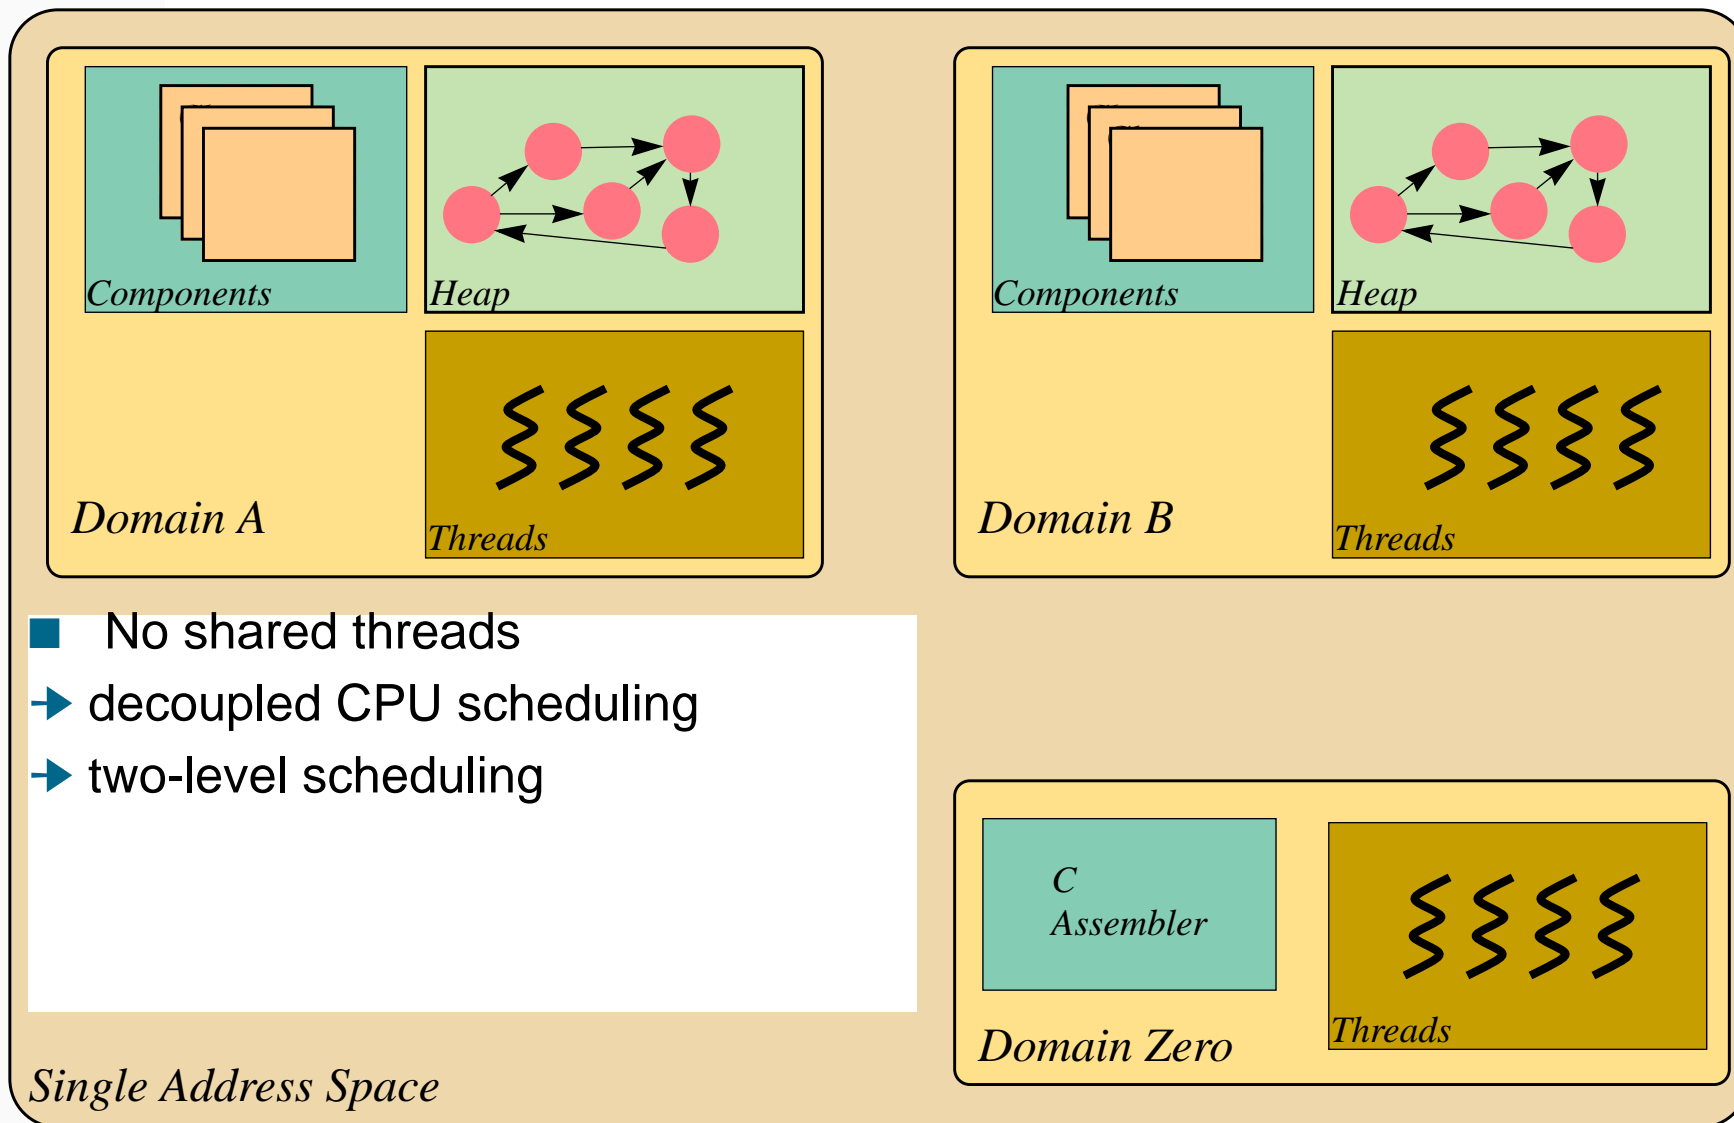
Objects & Heap

JX Architecture



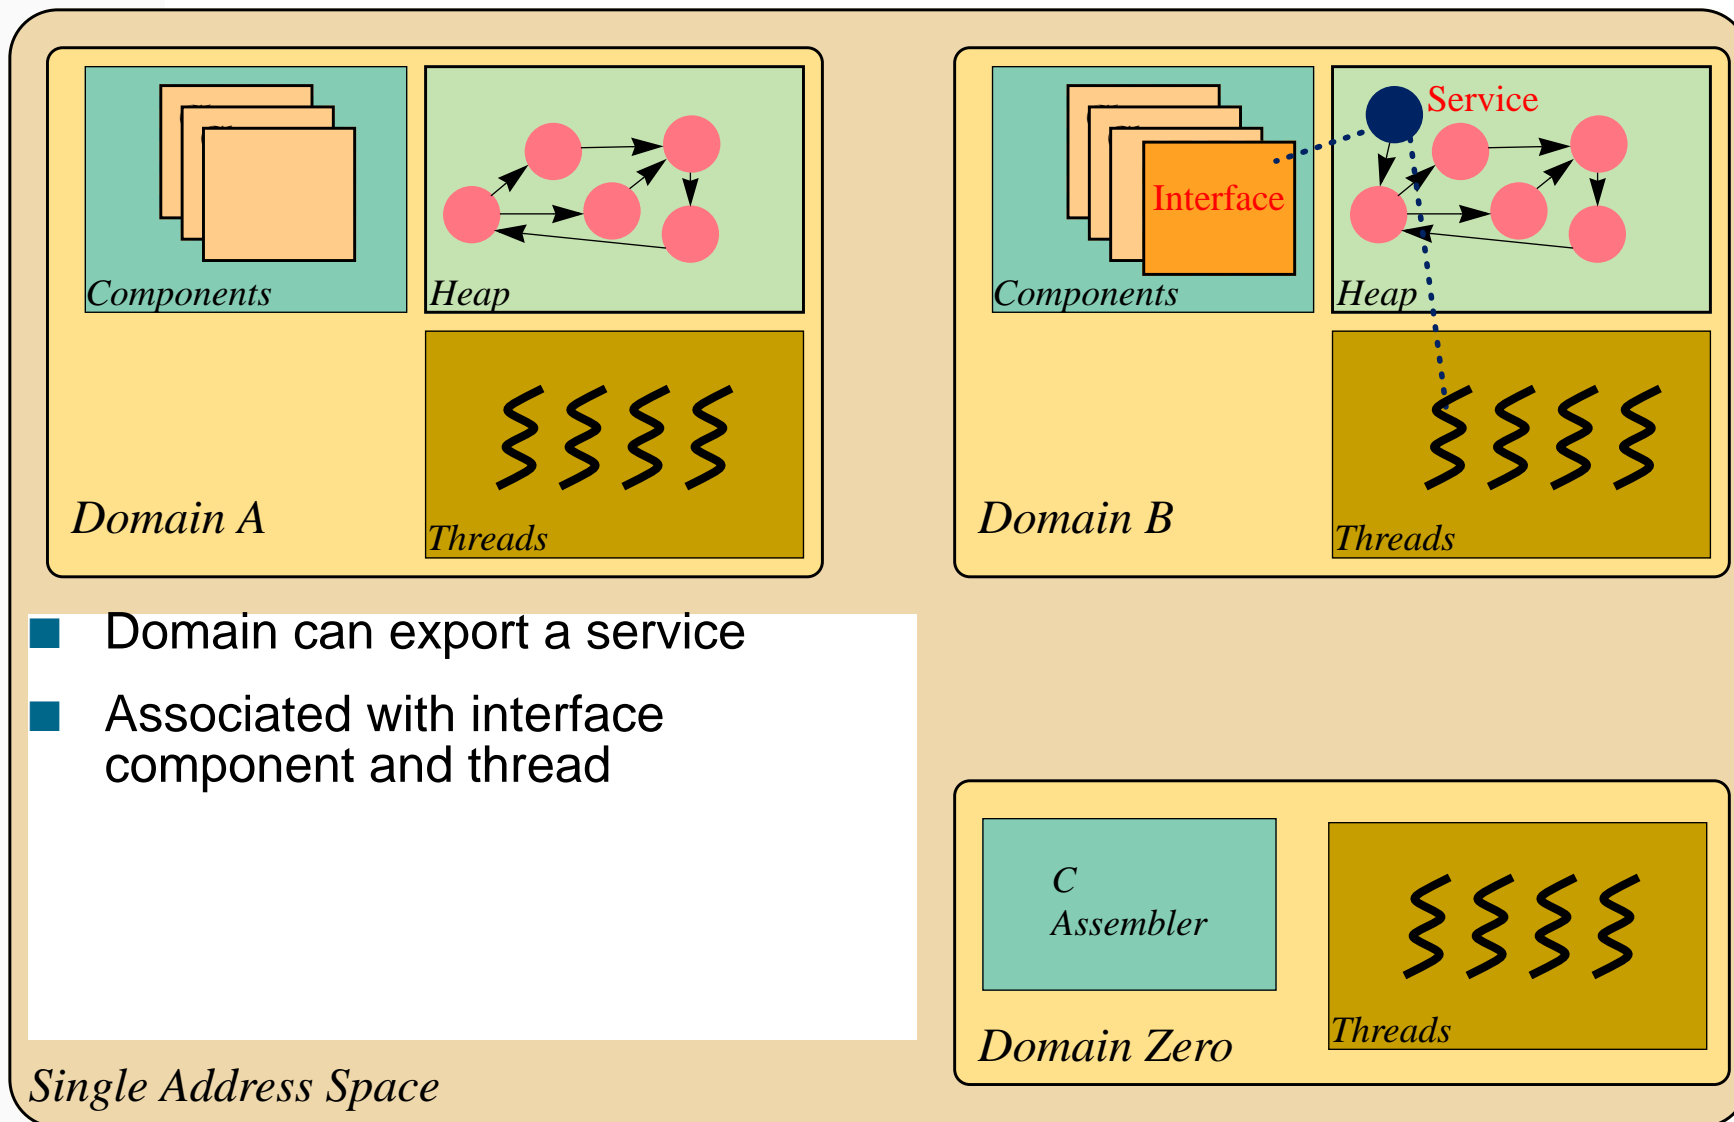
Threads

JX Architecture



Communication: Portals

JX Architecture

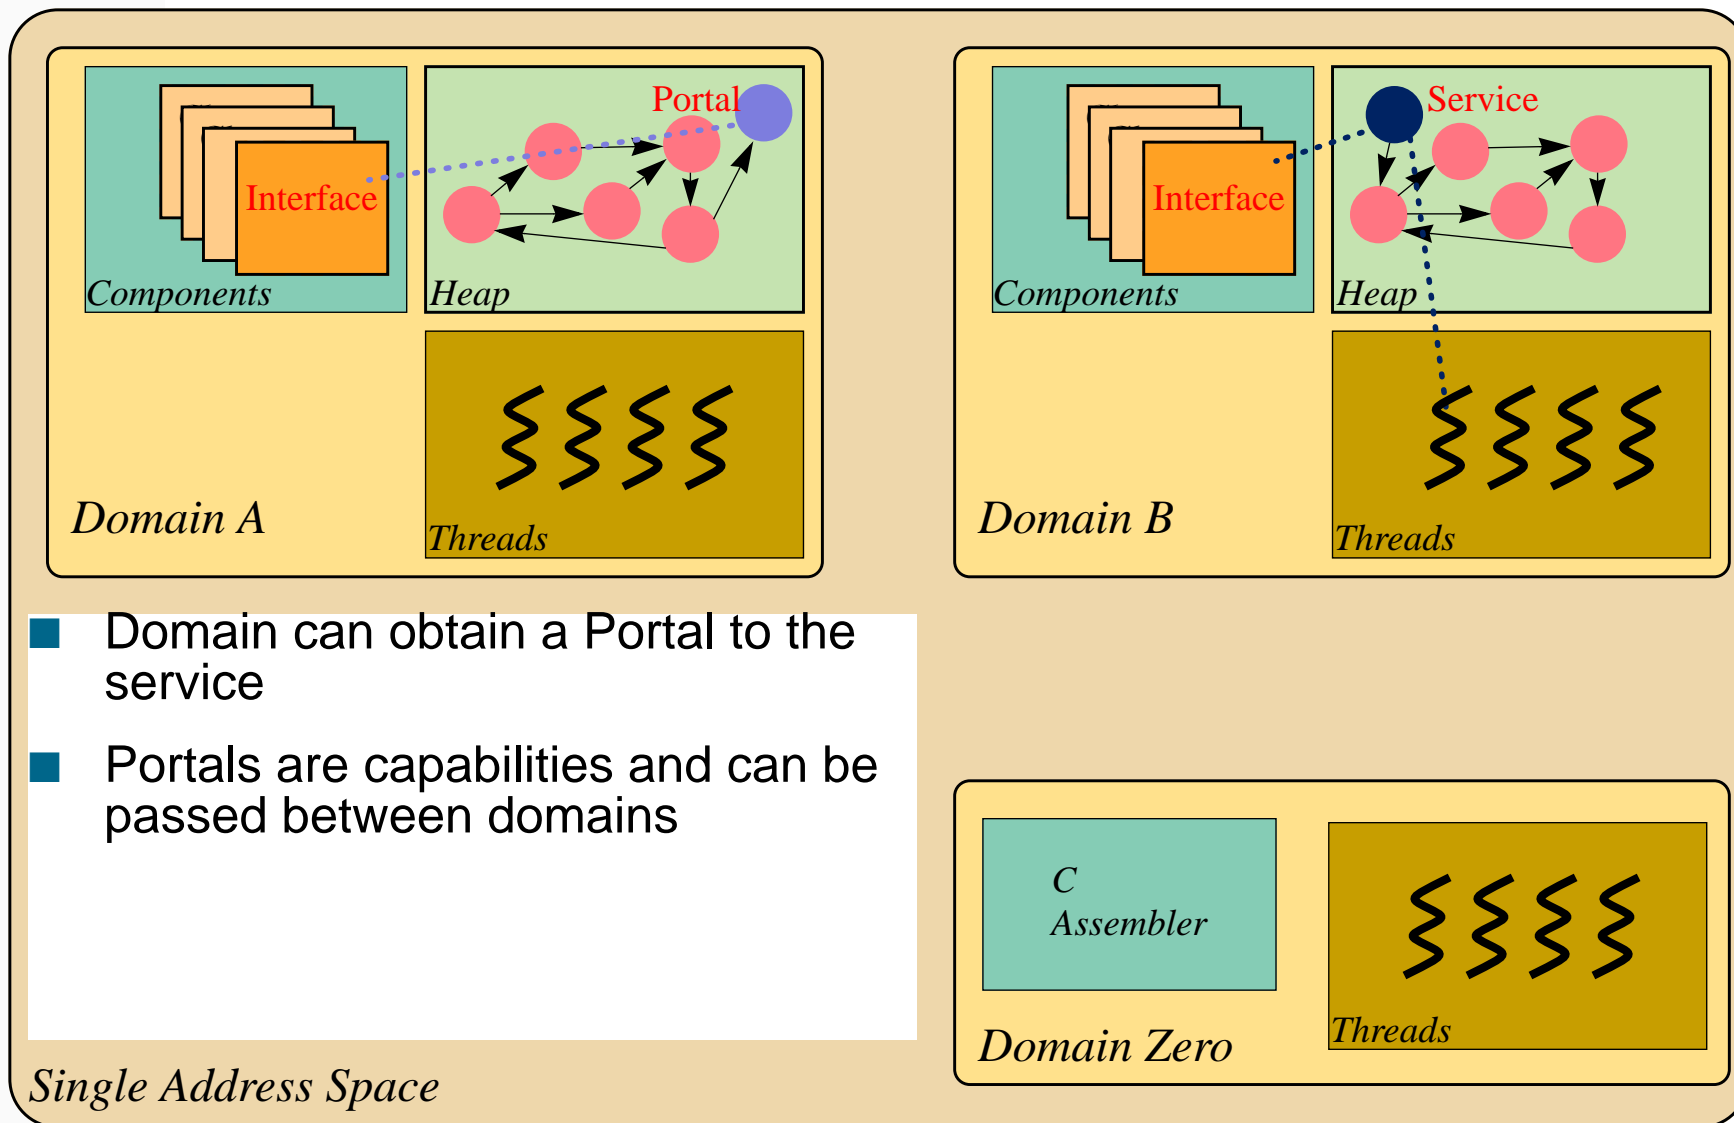


- Domain can export a service
- Associated with interface component and thread

Single Address Space

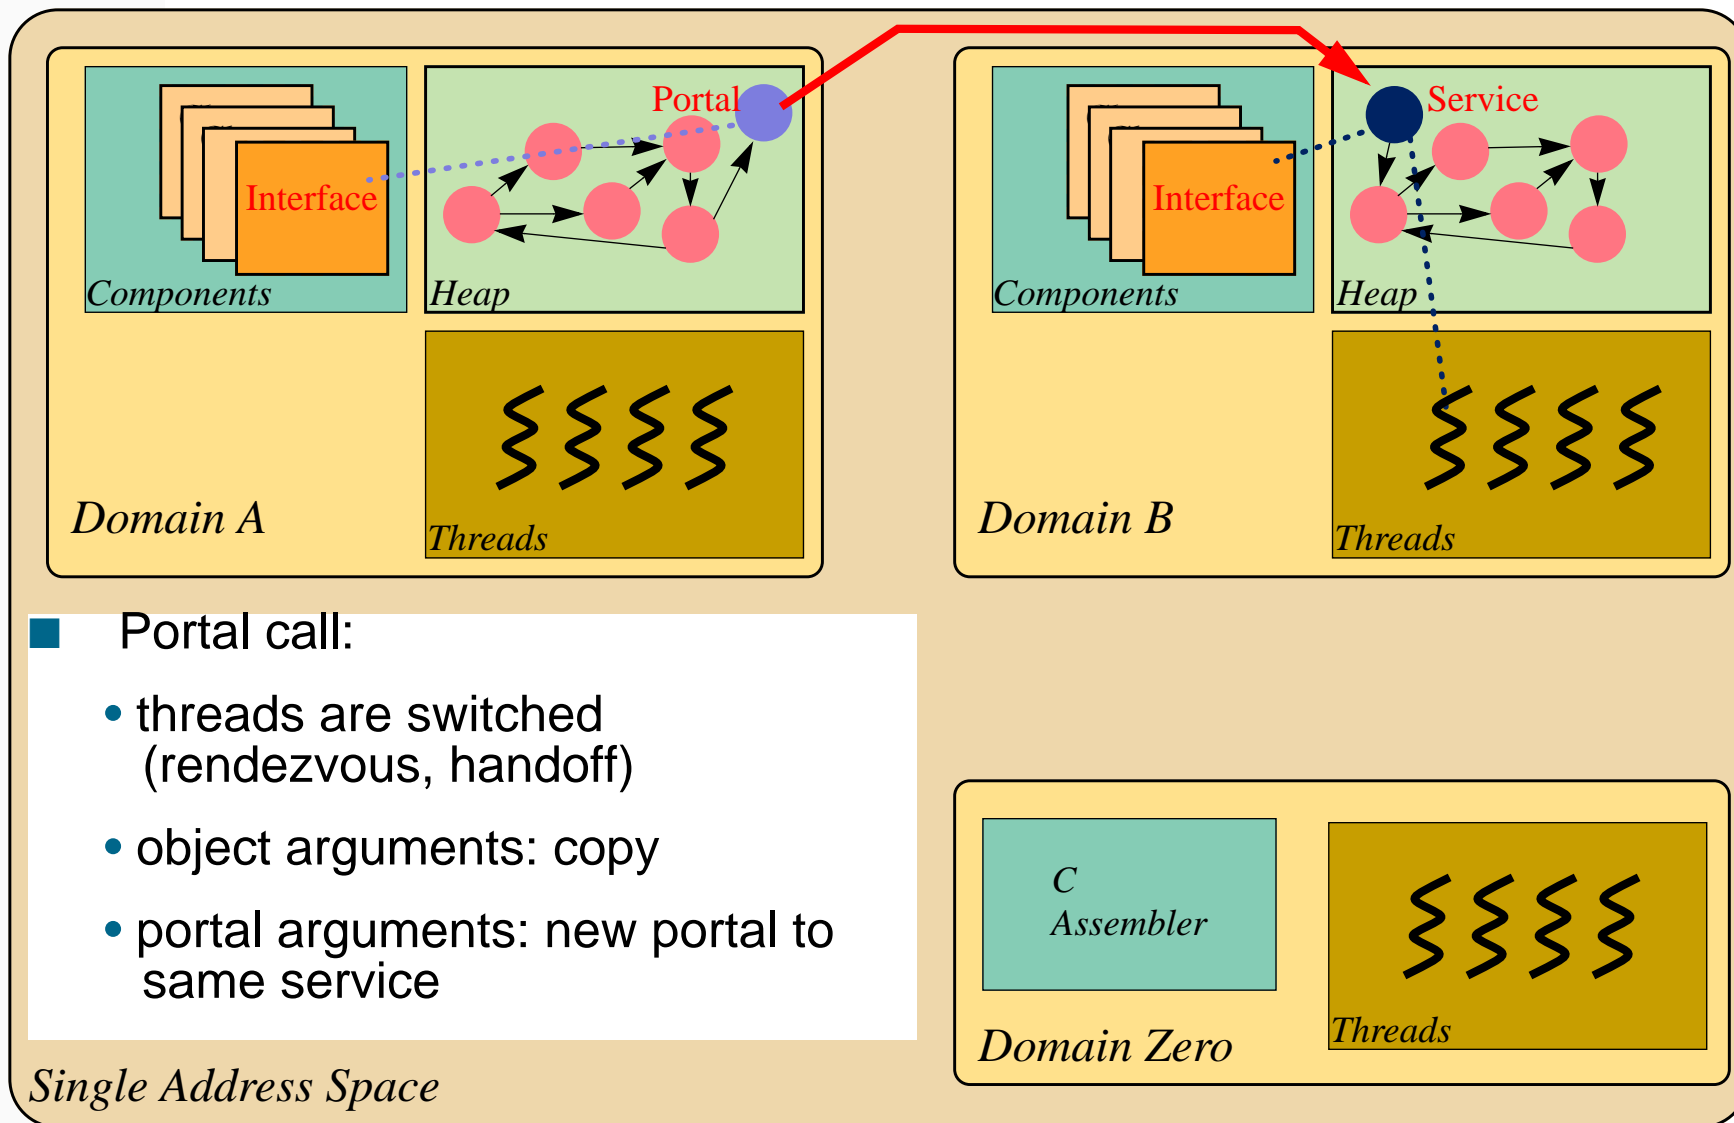
Communication: Portals

JX Architecture



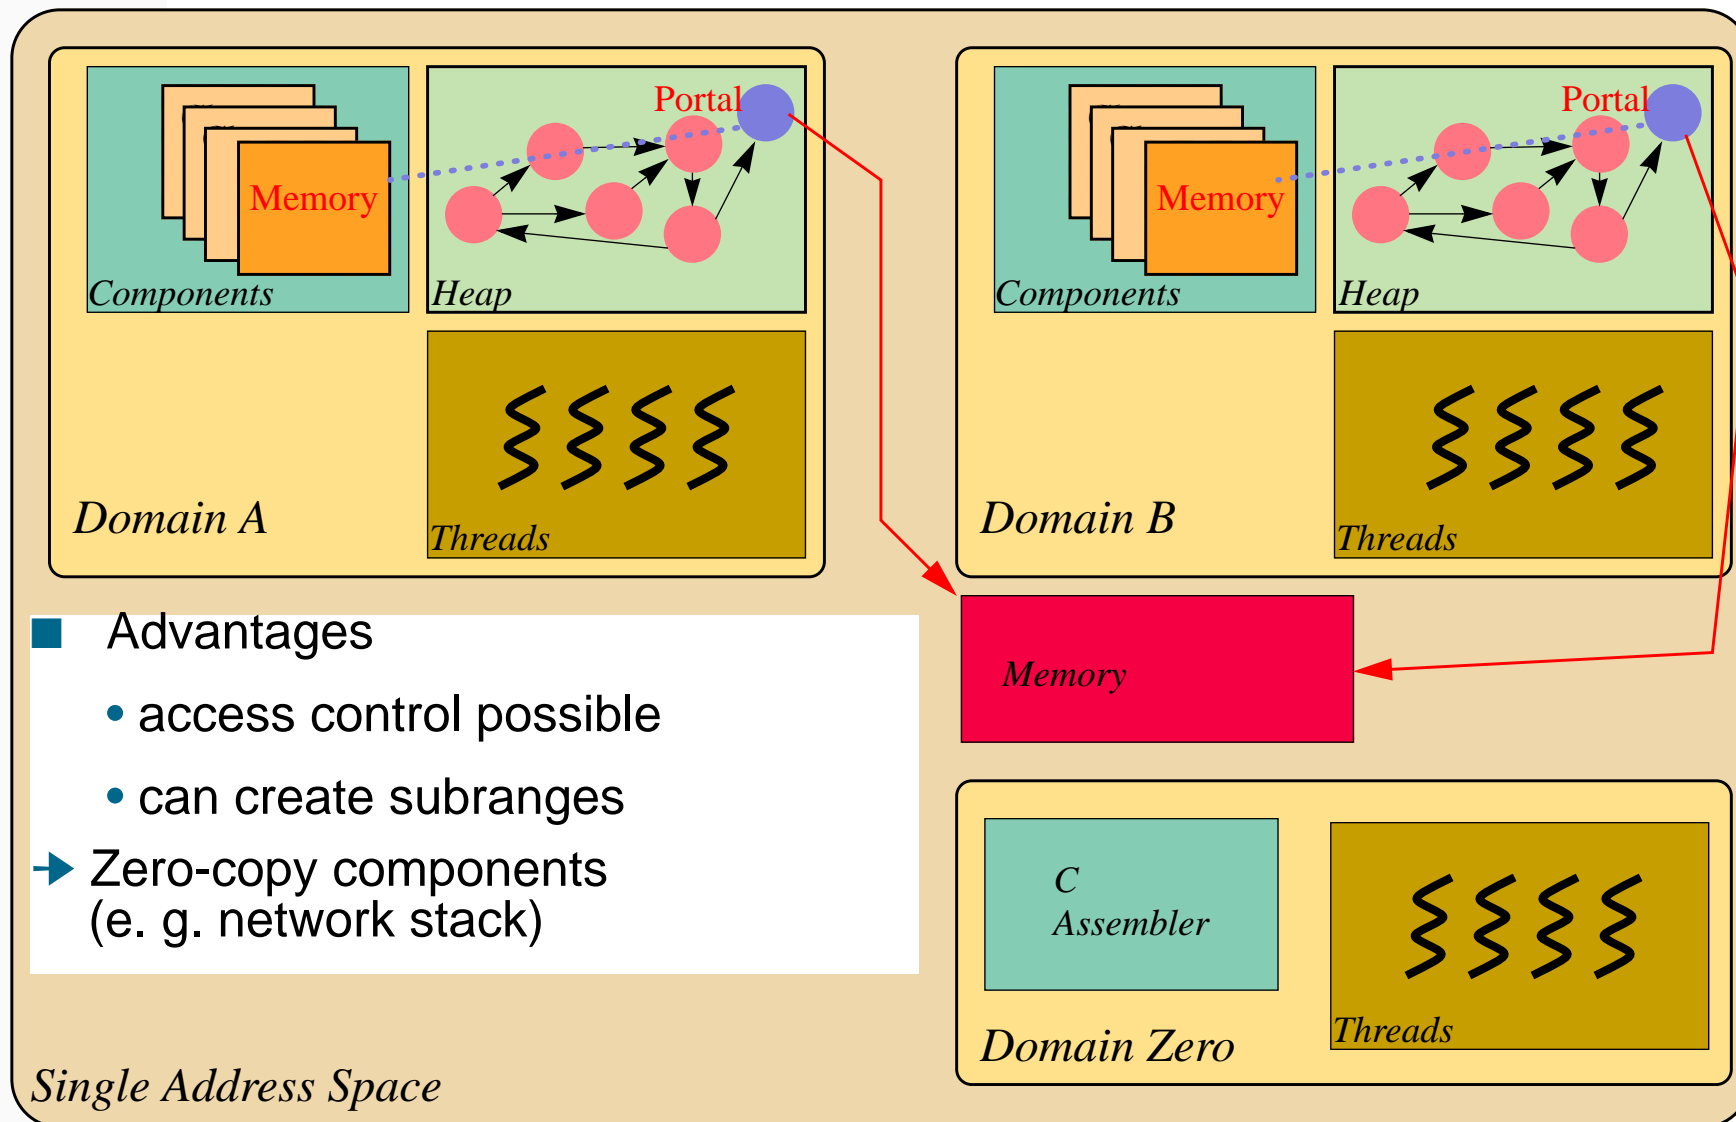
Communication: Portals

JX Architecture



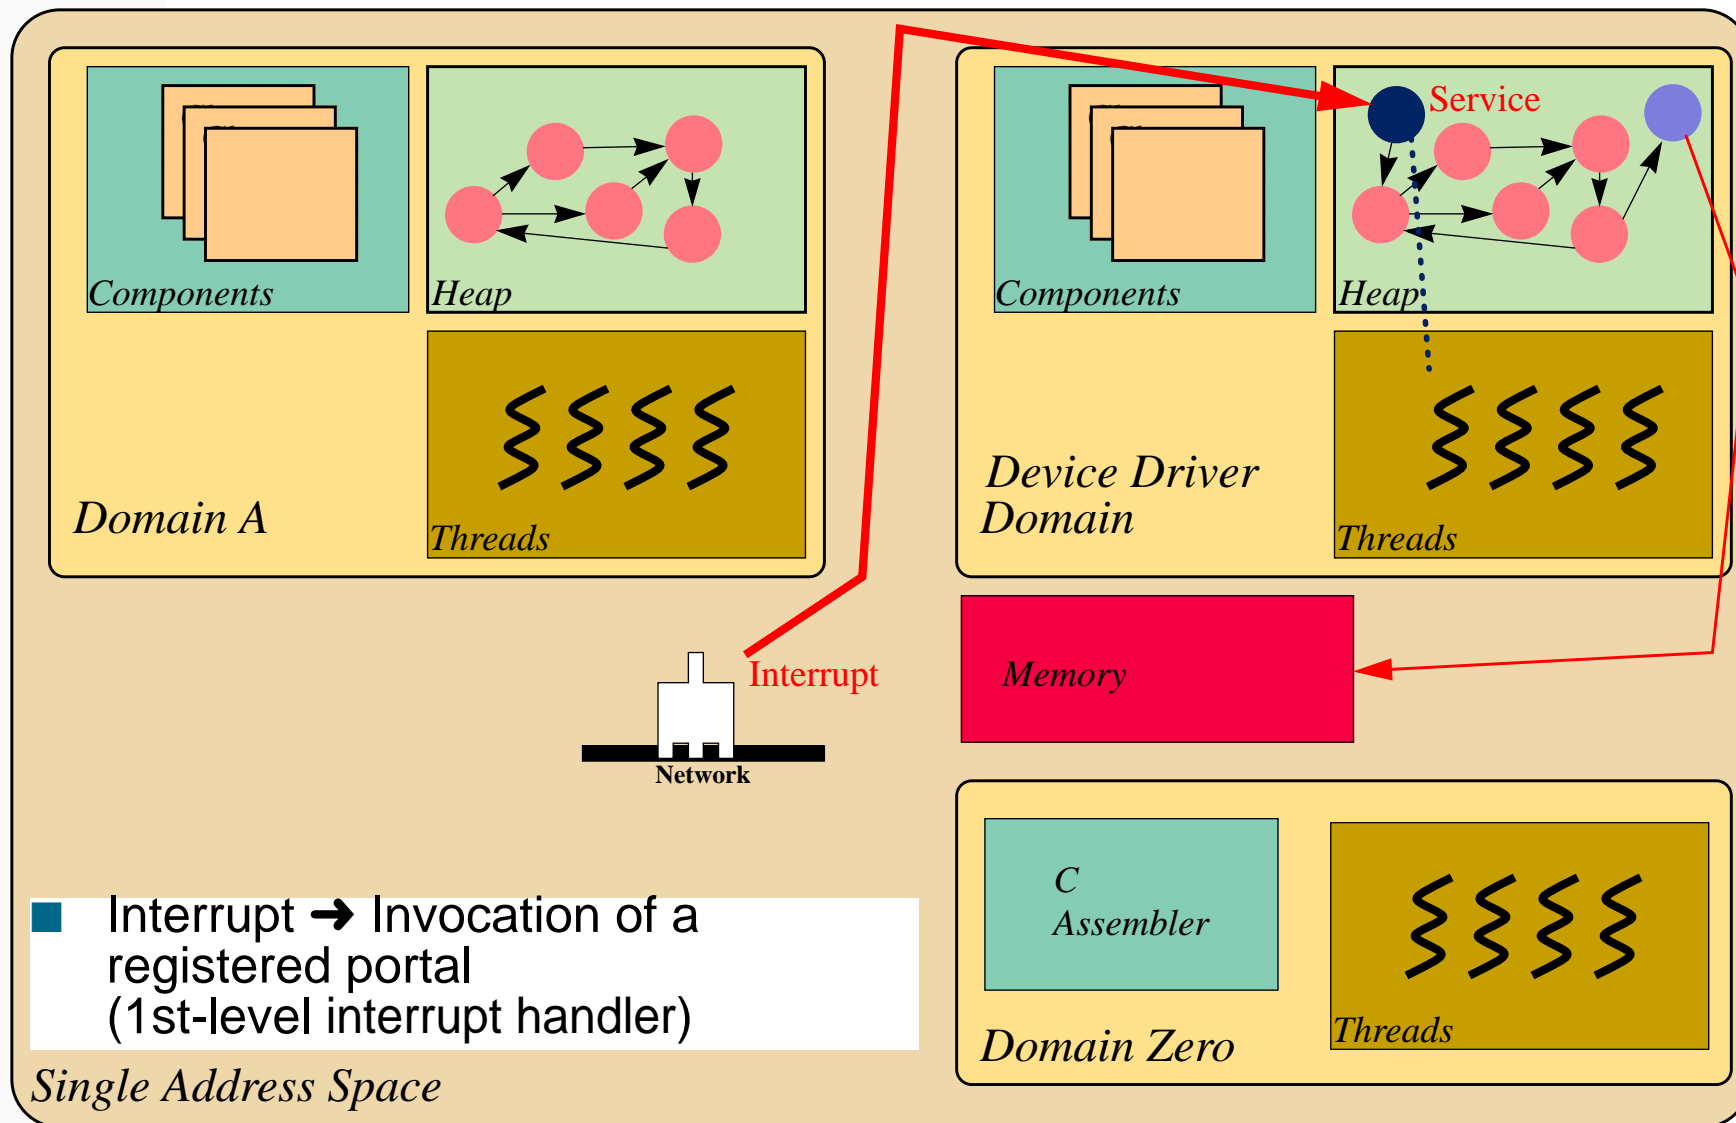
Communication: Memory

JX Architecture



Device Driver

JX Architecture

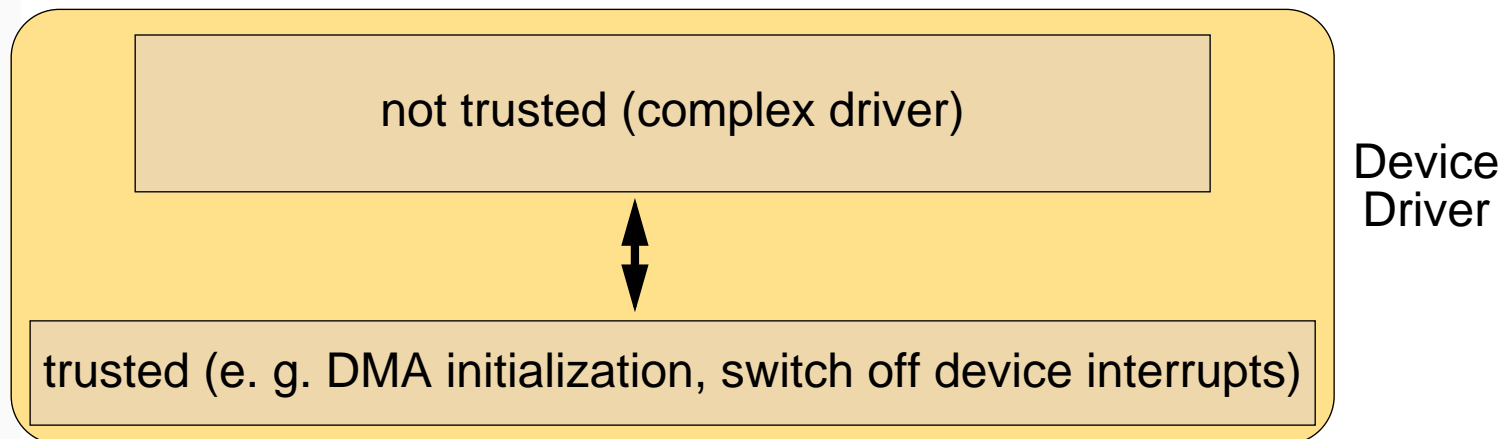


Protection: Interrupt Handler

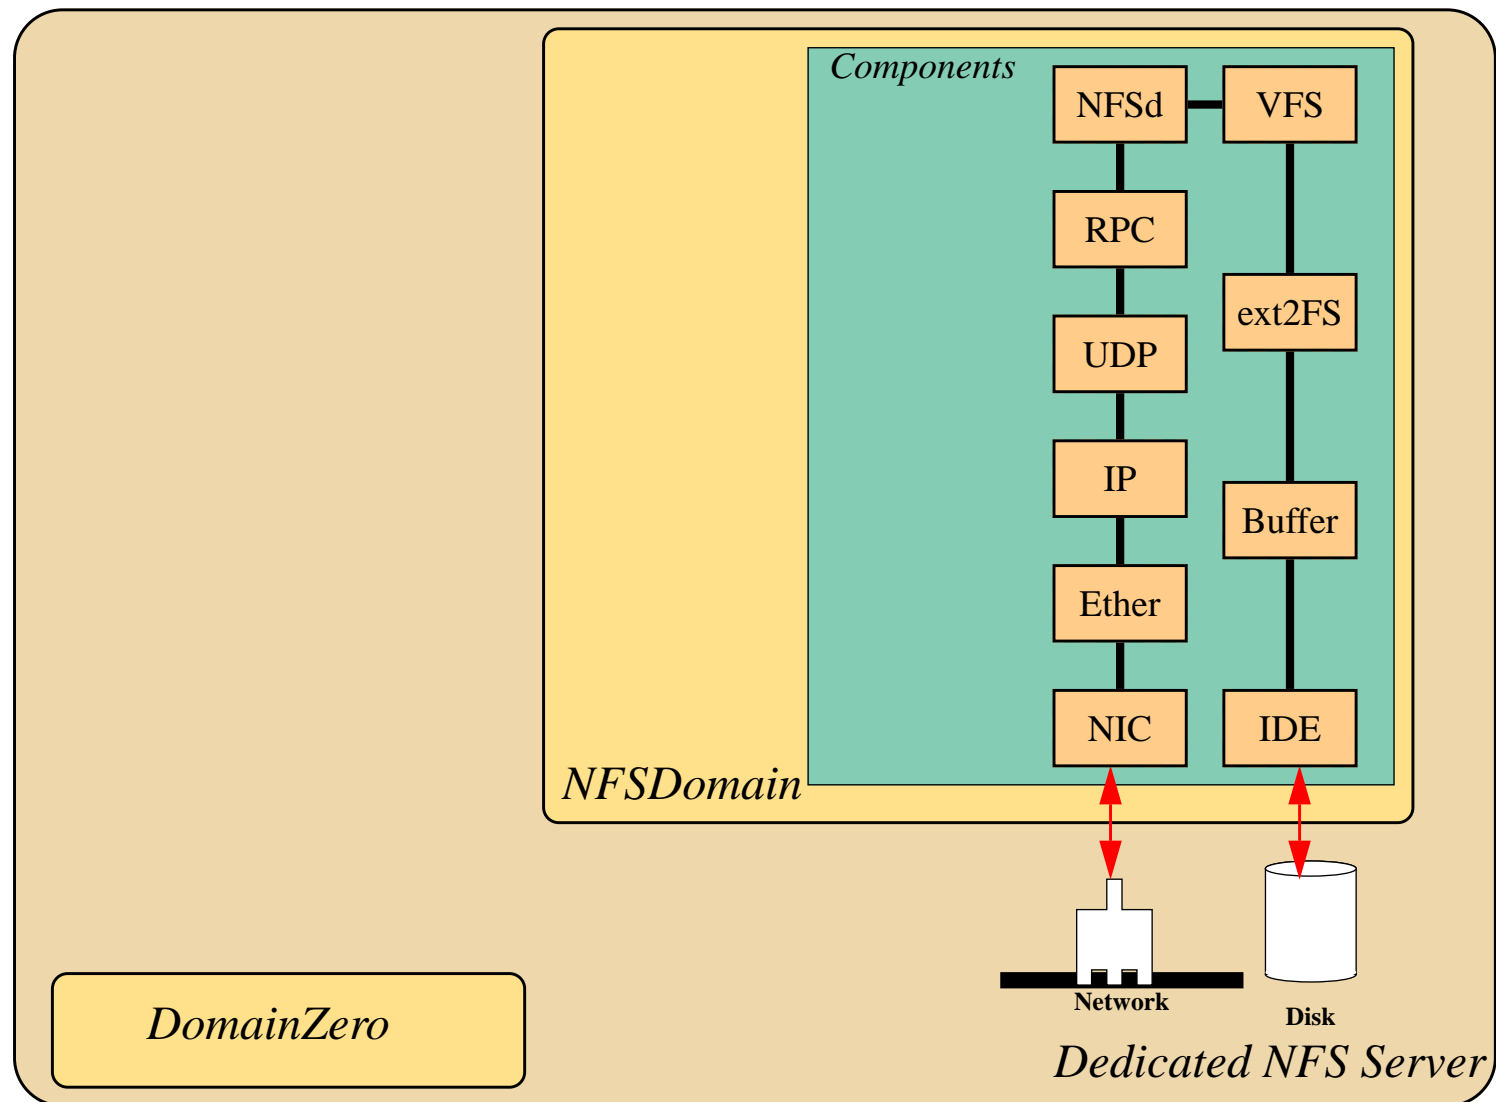
- Interrupts on the interrupted CPU are blocked during execution of the interrupt handler
- Verifier checks interrupt handler for upper limit of execution time
 - can insert runtime checks to ensure timely termination
 - runtime check can terminate interrupt handler and initiate counter measure (e. g., switch off device interrupts)

Protection: Device Driver

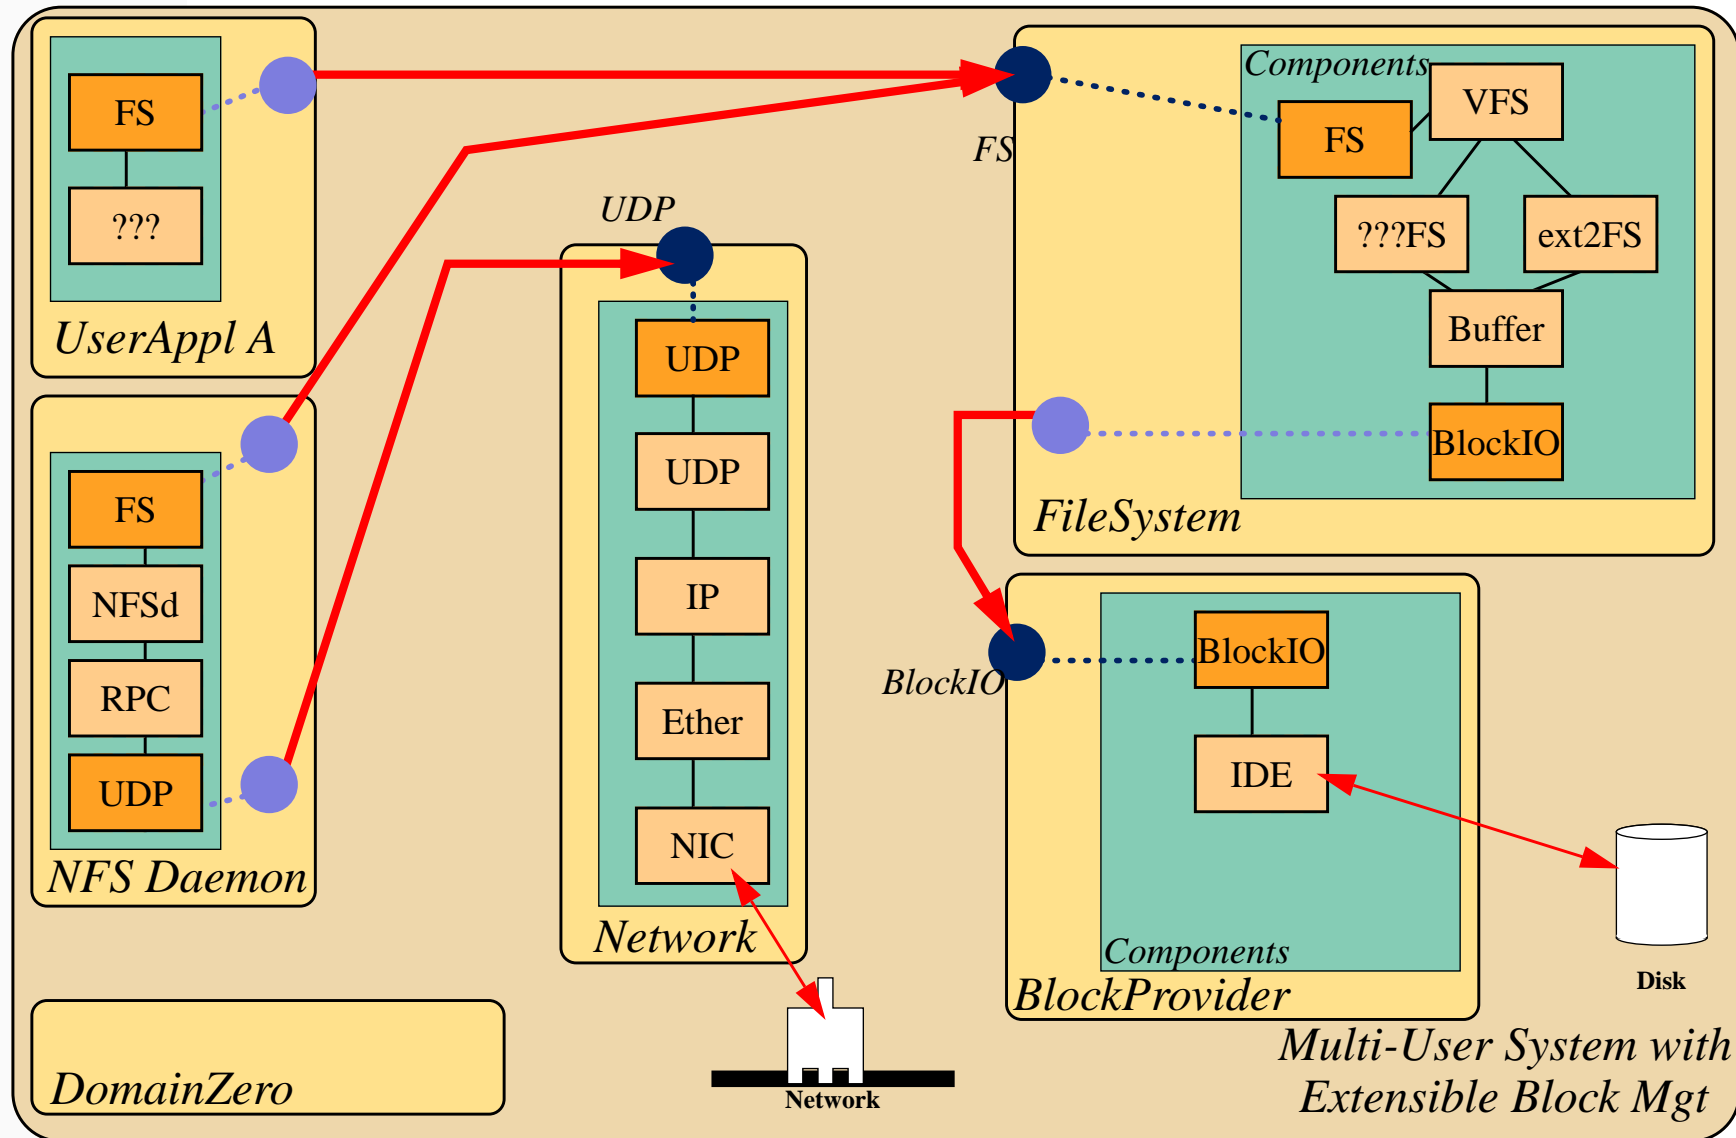
- JX protection is based on type safety and portals
- Some domains can circumvent these mechanisms
 - DomainZero, Translator, Verifier → Trust
 - (some) device drivers



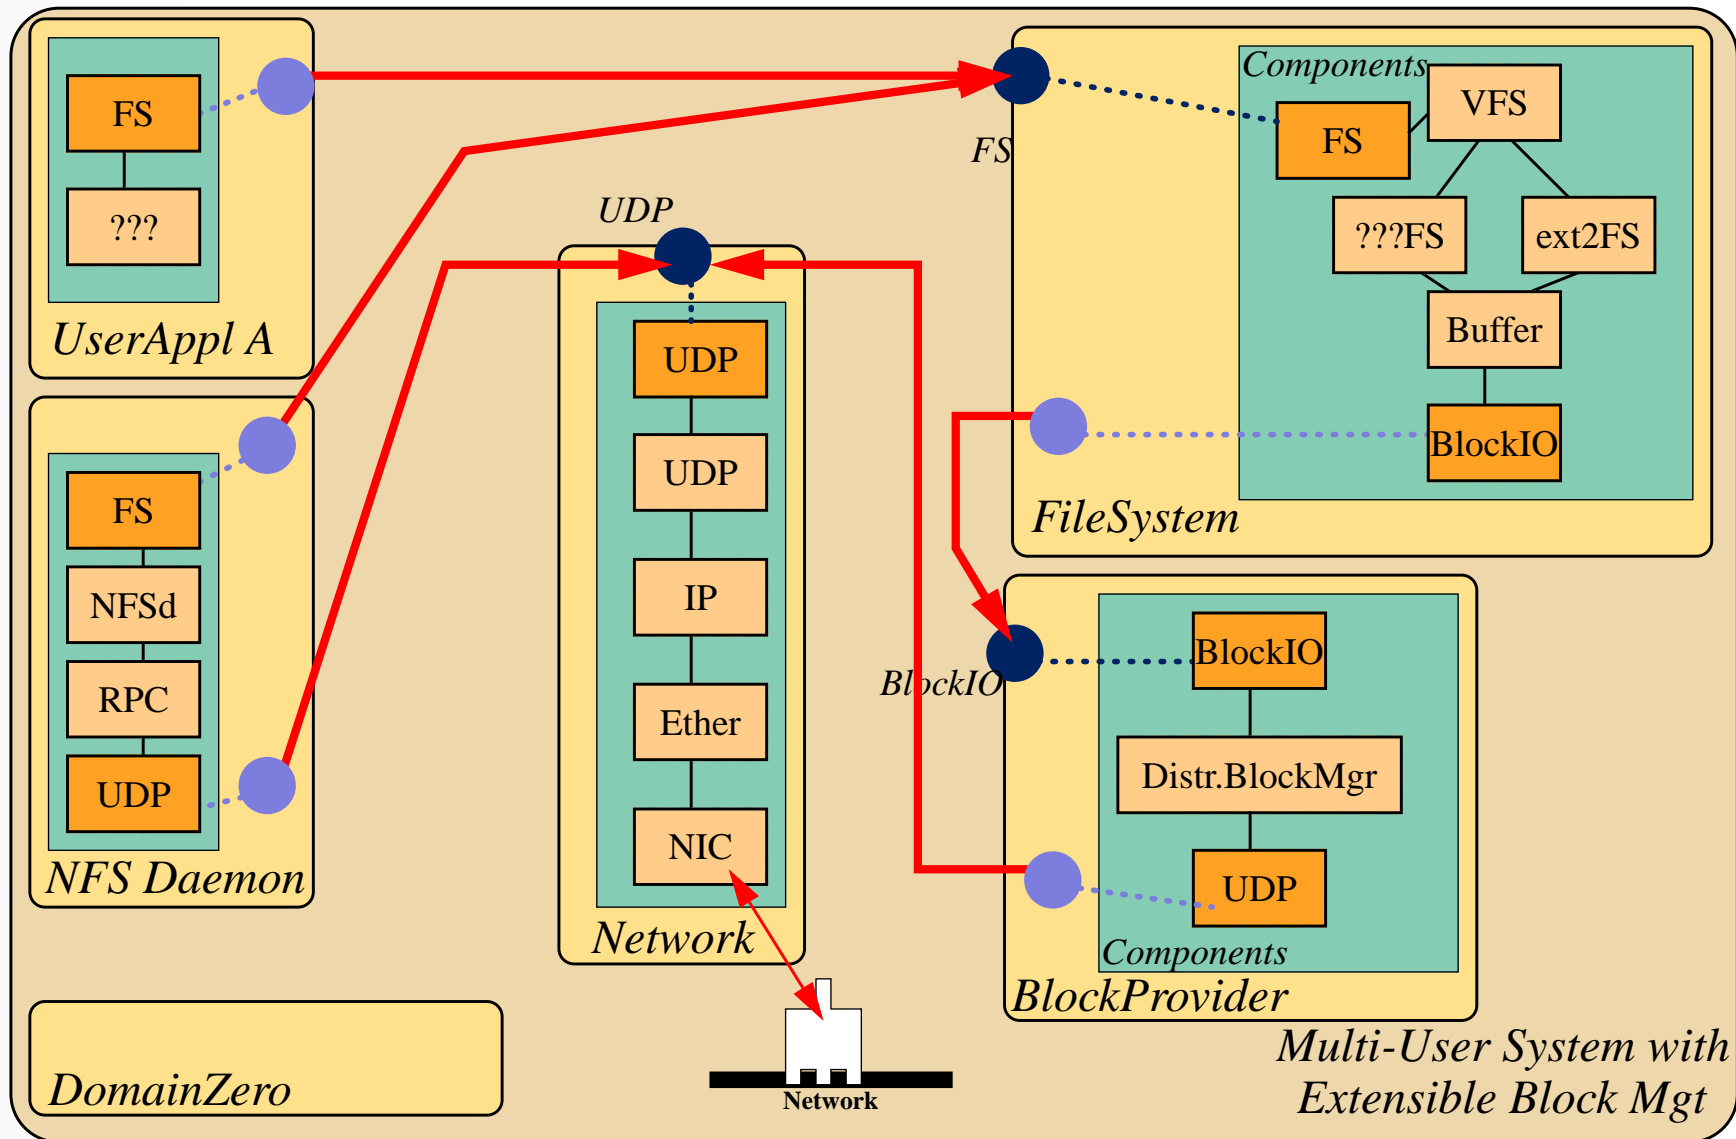
Building an OS: A Dedicated System



Building an OS: Extensibility



Building an OS: Extensibility




Performance

- IPC:
 - Portal round trip 650 cycles
 - L4 (including RPC stubs): 800 cycles
 - KaffeOS: 27270 cycles


Performance

- lozone-like benchmark: record=4kB, file=4kB, re-read
 - JX (PIII 500): 201 MB/s
 - Linux (PIII 500): 400 MB/s



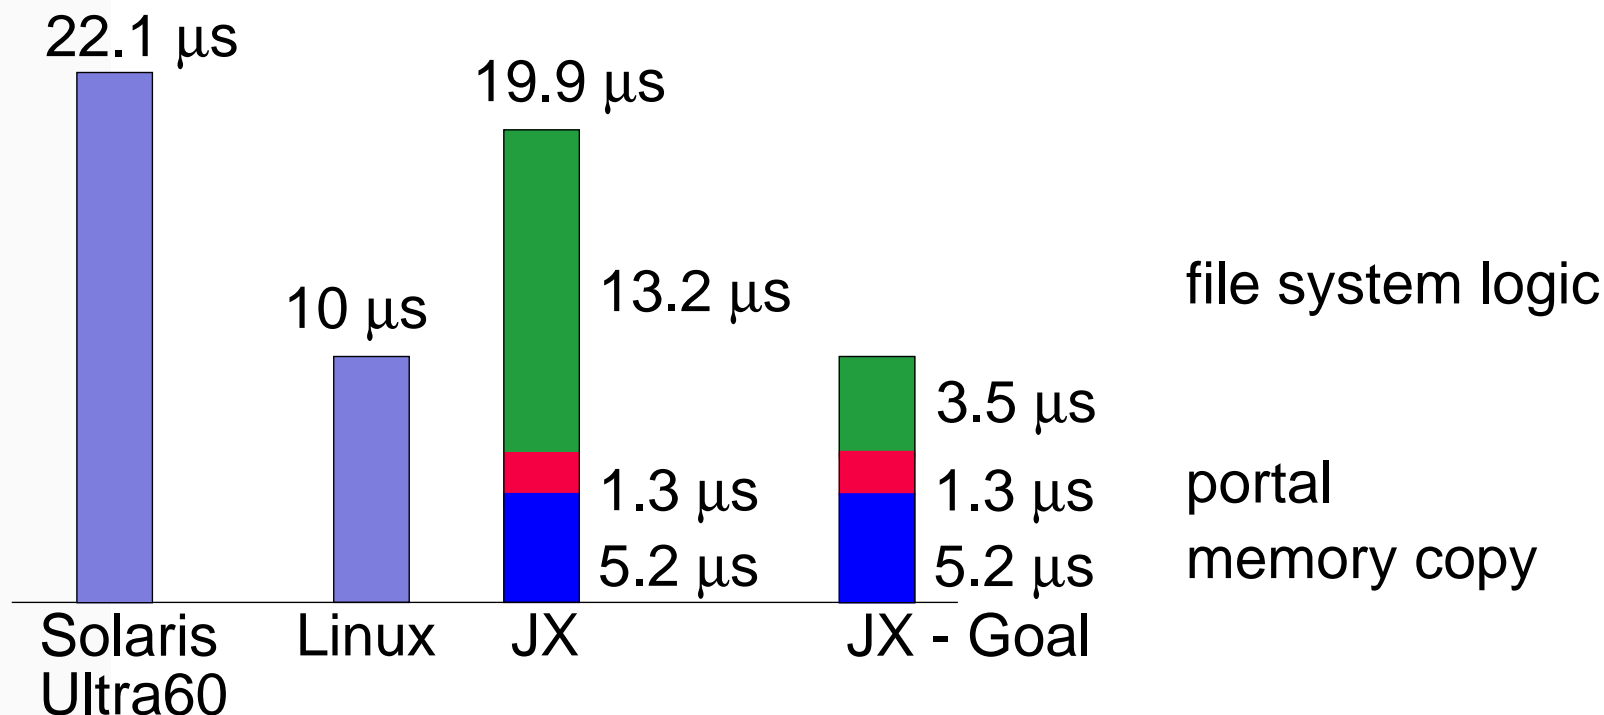
Performance

- lozone-like benchmark: record=4kB, file=4kB, re-read
 - JX (PIII 500): 201 MB/s
 - Linux (PIII 500): 400 MB/s



Performance

- lozone-like benchmark: record=4kB, file=4kB, re-read
 - JX (PIII 500): 201 MB/s
 - Linux (PIII 500): 400 MB/s
 - Solaris (Ultra60 - UltraSparcII 360): 181 MB/s



Conclusion

- Single Address Space
- Full Protection
 - completely decoupled domains
 - fast communication using portals or memory objects
- Reusable components
- Dynamic extensibility
- Good performance

→ <http://www4.cs.fau.de/Projects/JX/>

Status

- JX runs on off-the-shelf PCs
- Drivers for: IDE, Matrox G200, 3COM 3C905B, BT848
- Ext2-FS
- UDP, TCP, RPC (client), NFS (client)
- SMP support