

The JX Operating System: An Overview

Michael Golm



University of Erlangen
Department of Computer Science
(Operating Systems and Distributed Systems)
Martensstrasse 1
91058 Erlangen, Germany
golm@cs.fau.de

Outline

- Objectives
- Architecture
- Protection
- Flexibility
- Performance

Objectives of the JX system

- Make writing an OS as easy as writing applications
 - ← simple and robust architecture
- Dynamic OS extension with untrusted components
 - ← exact resource accounting and customizable management
- Code reuse
 - tailored OS configurations; dedicated systems
- Protection
 - flexibility
 - performance
 - robustness

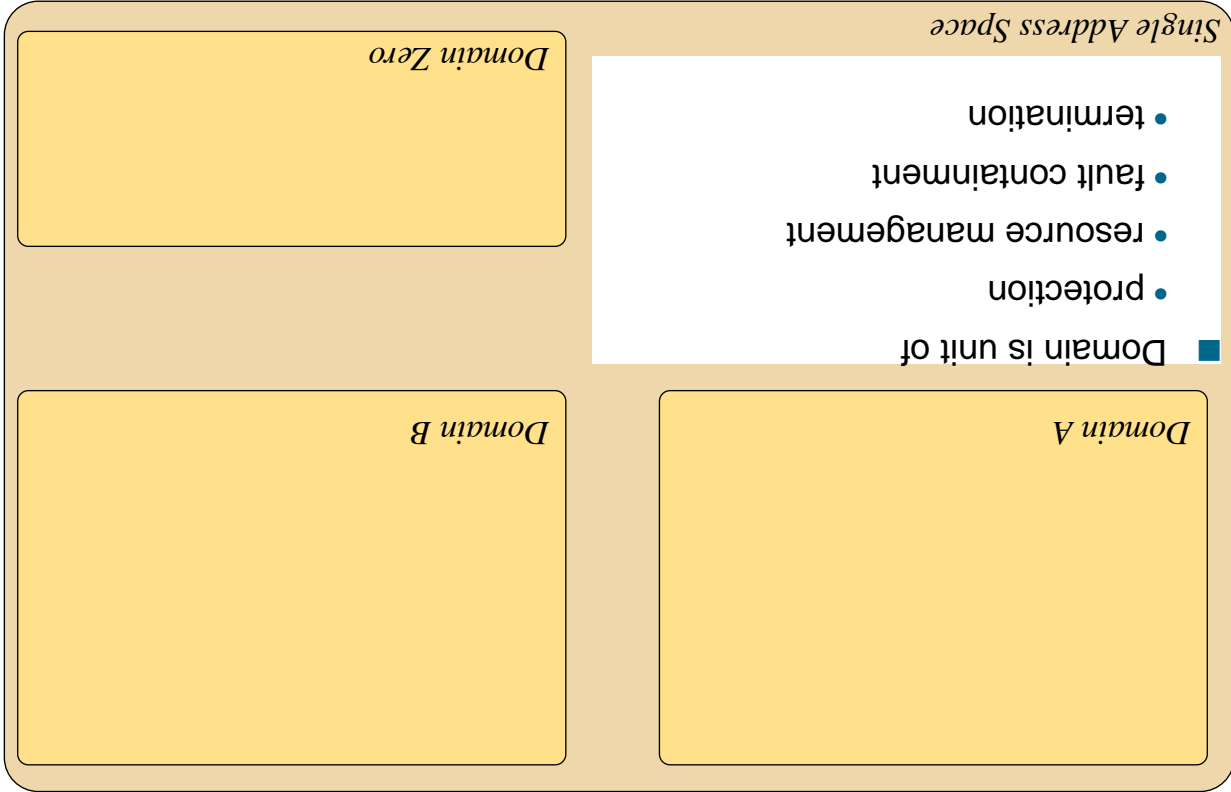
Single Address Space

JX Architecture

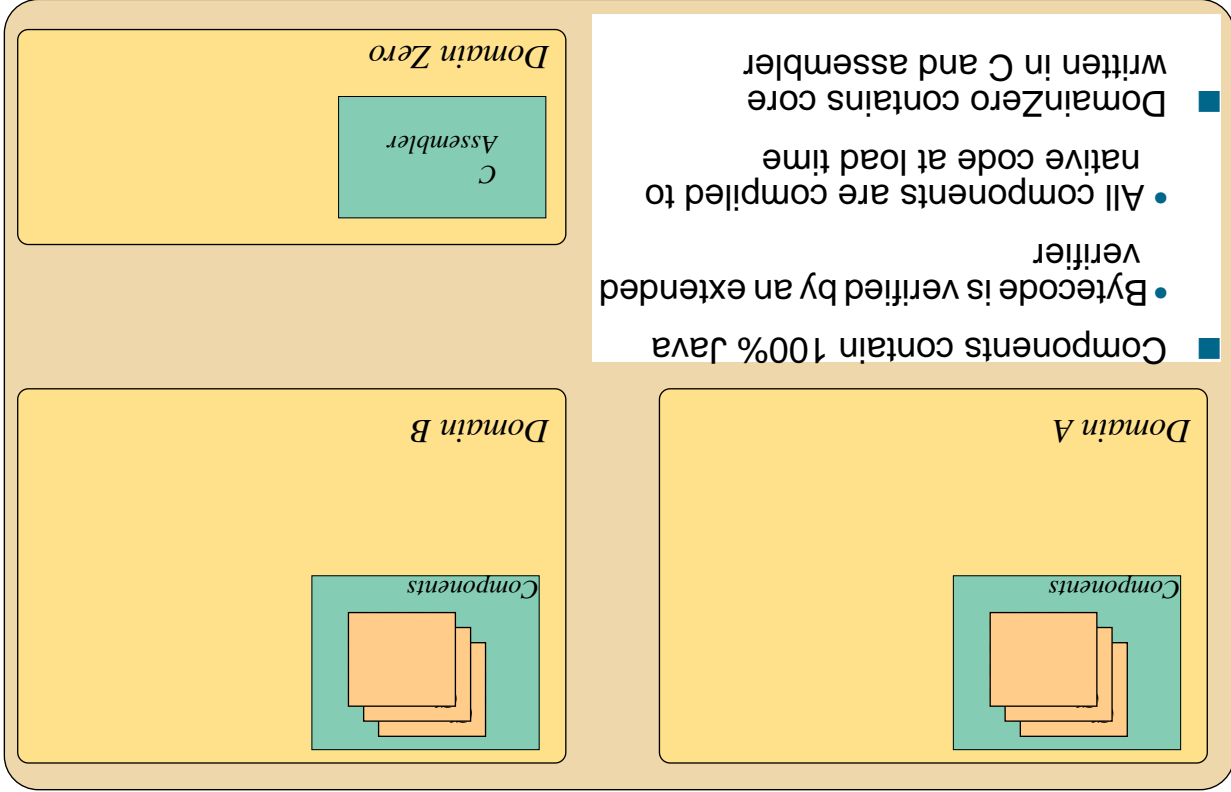
Single Address Space



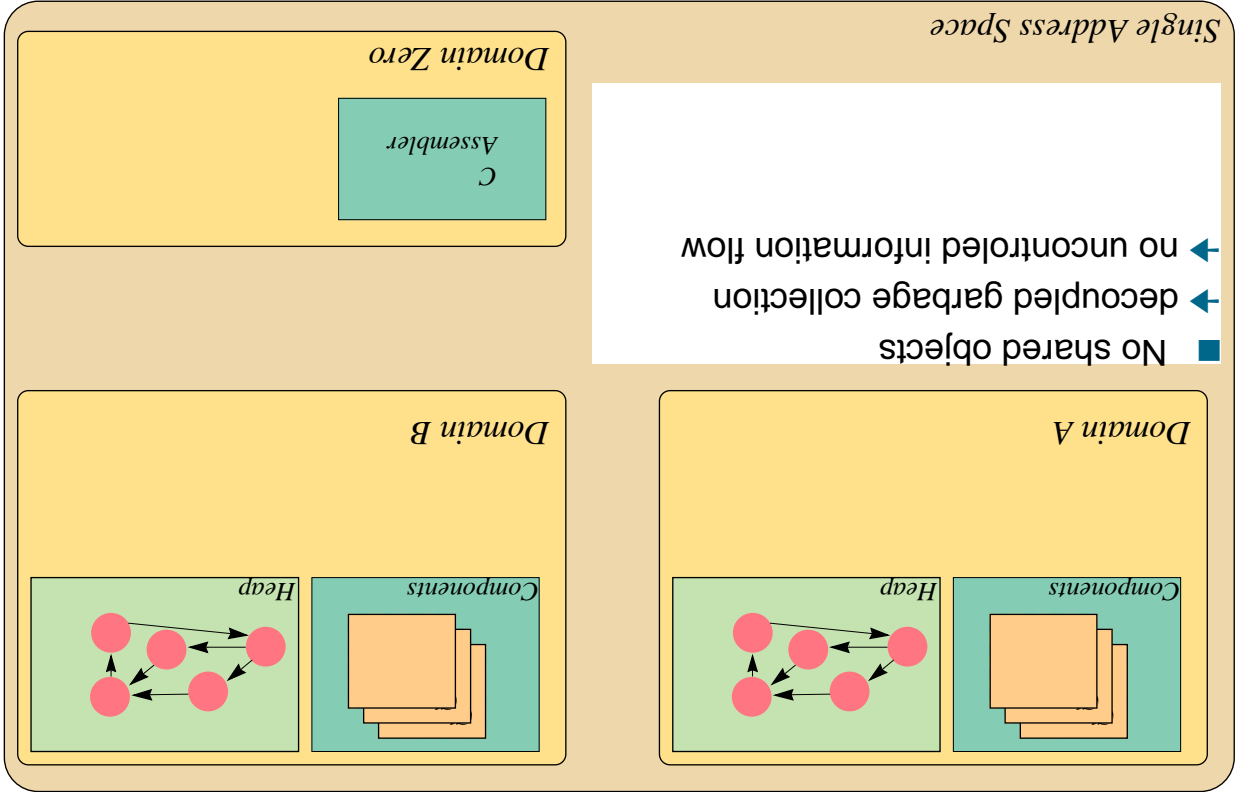
Domains



Components

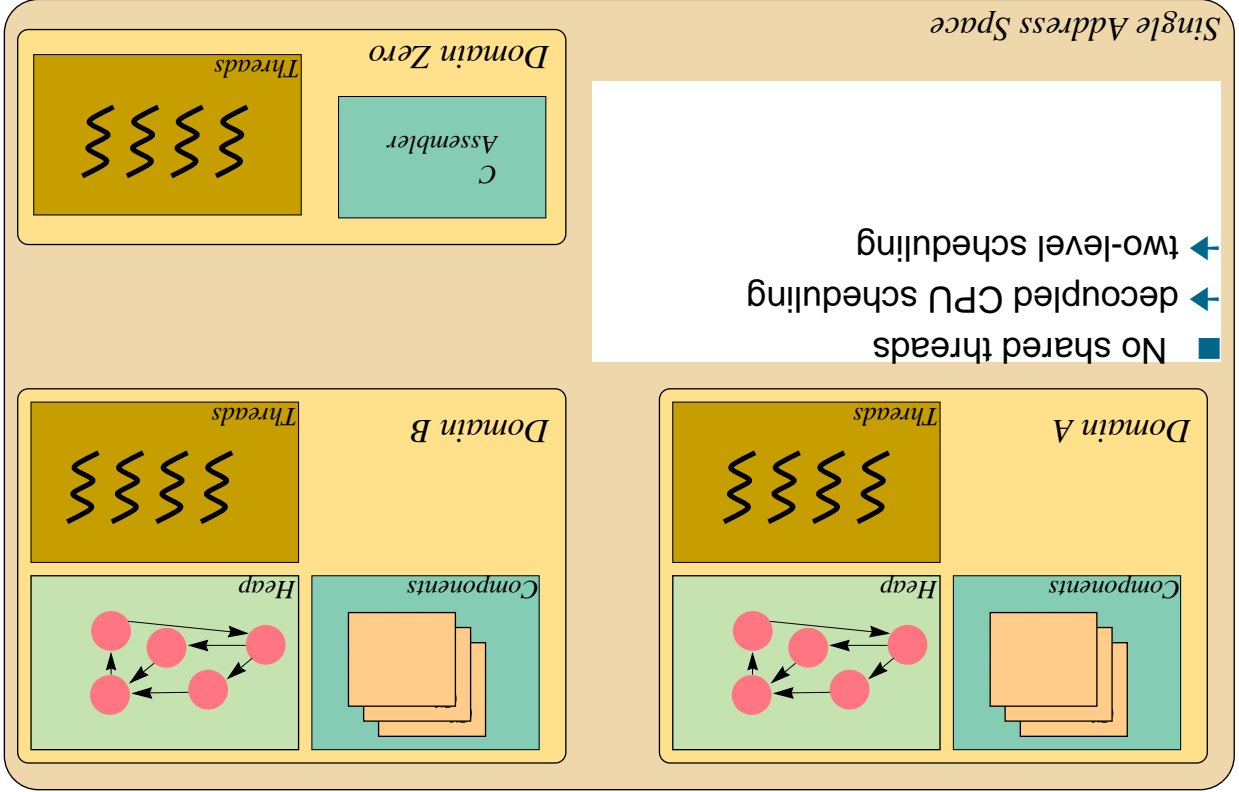


Objects & Heap



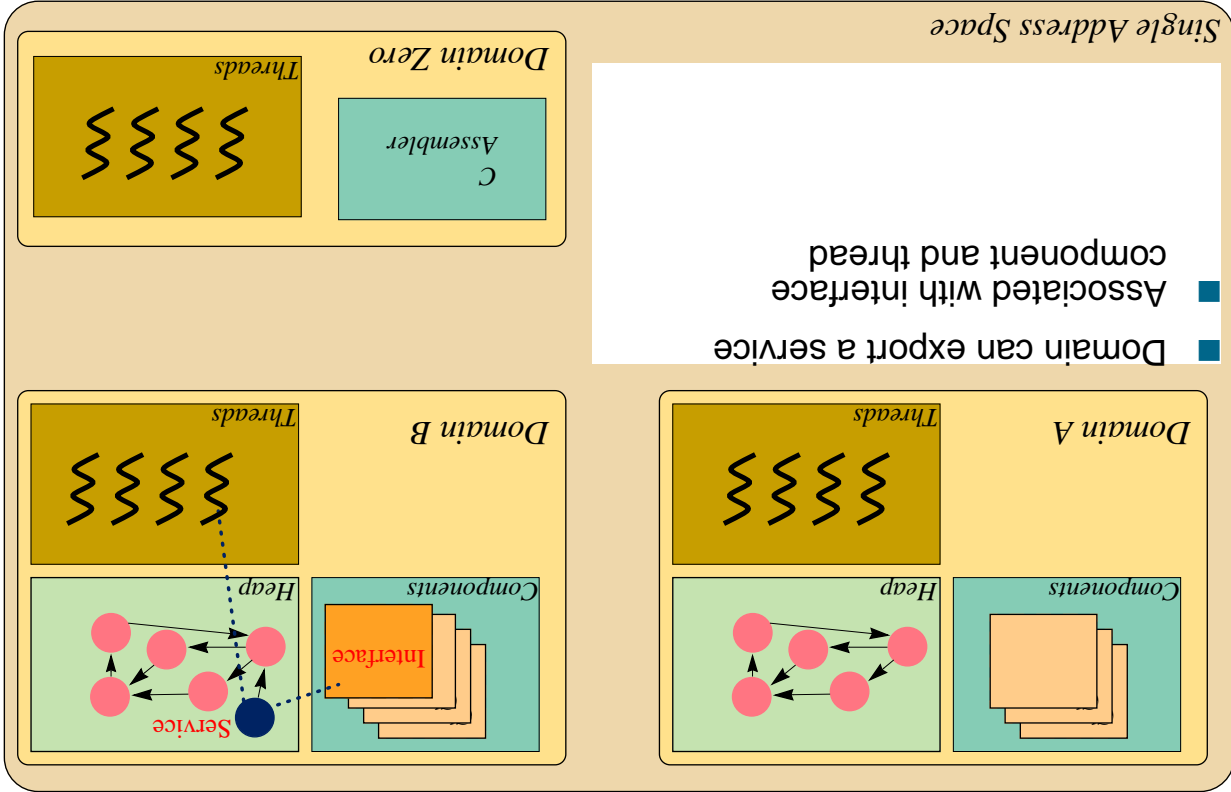
- No shared objects
- ▶ decoupled garbage collection
- ▶ no uncontrolled information flow

Threads

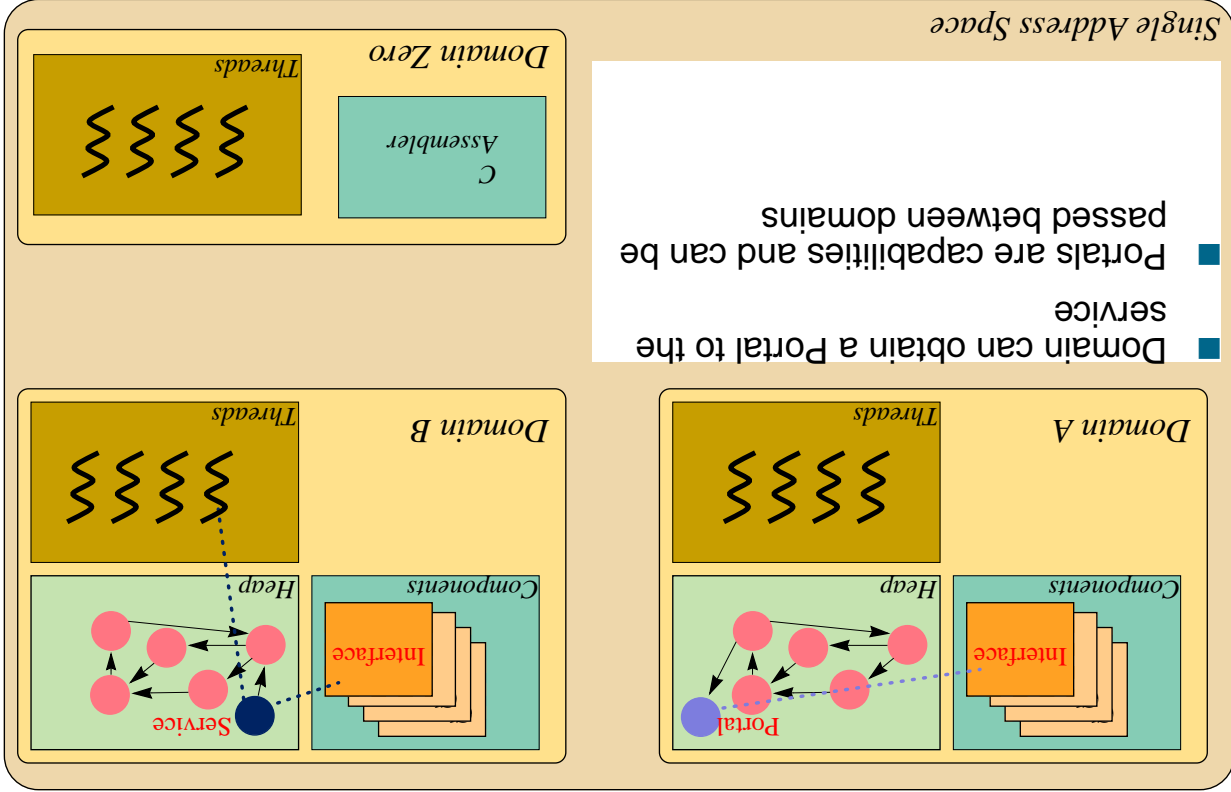


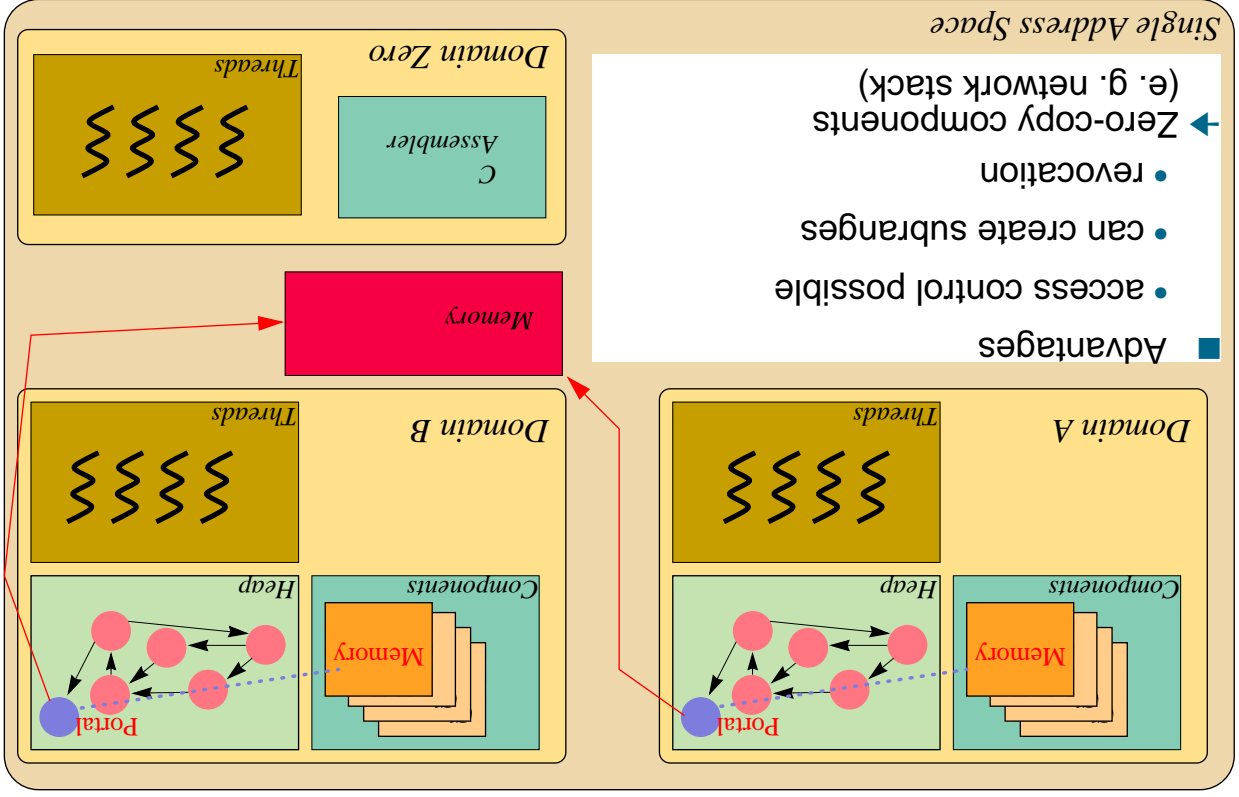
- No shared threads
- ▶ decoupled CPU scheduling
- ▶ two-level scheduling

Communication: Portals

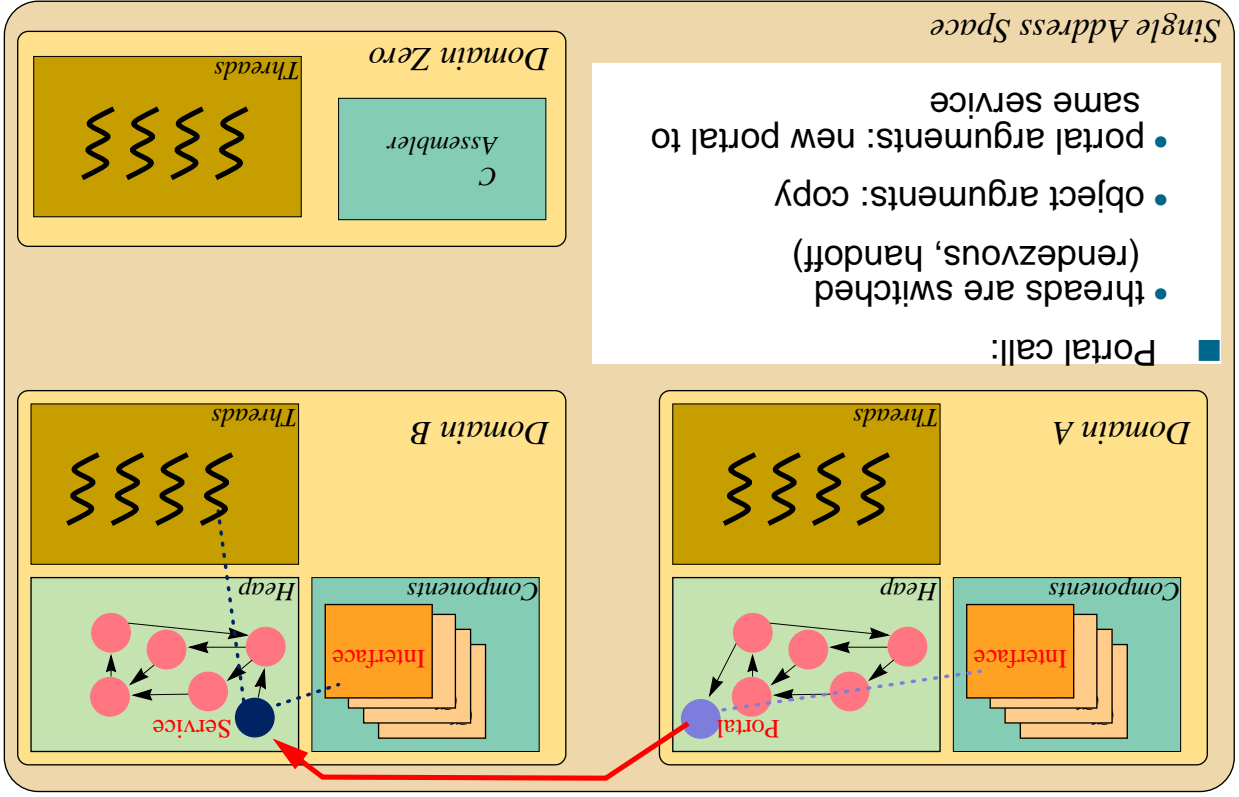


Communication: Portals

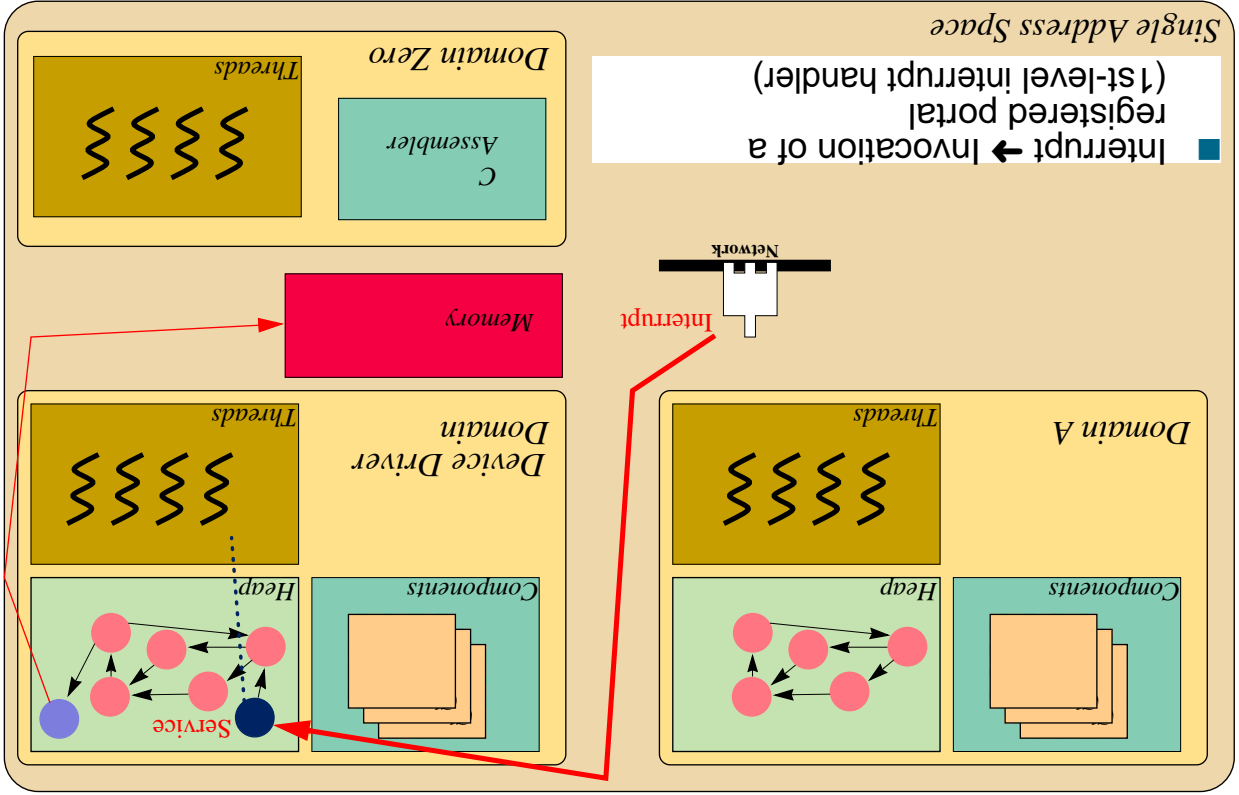




Communication: Memory



Communication: Portals

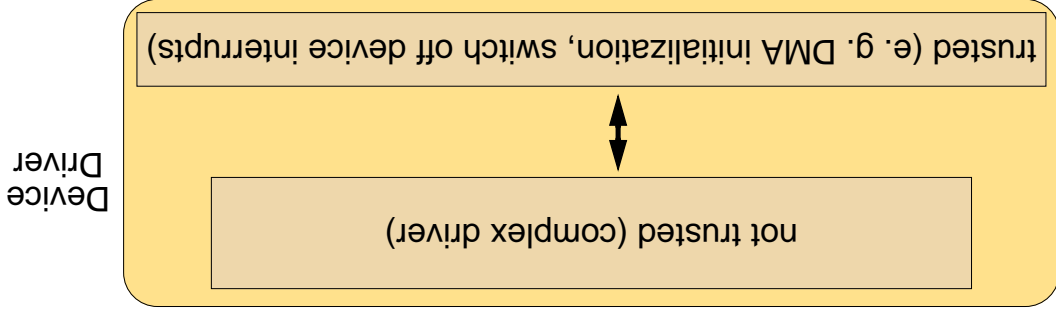


Protection: Interrupt Handler

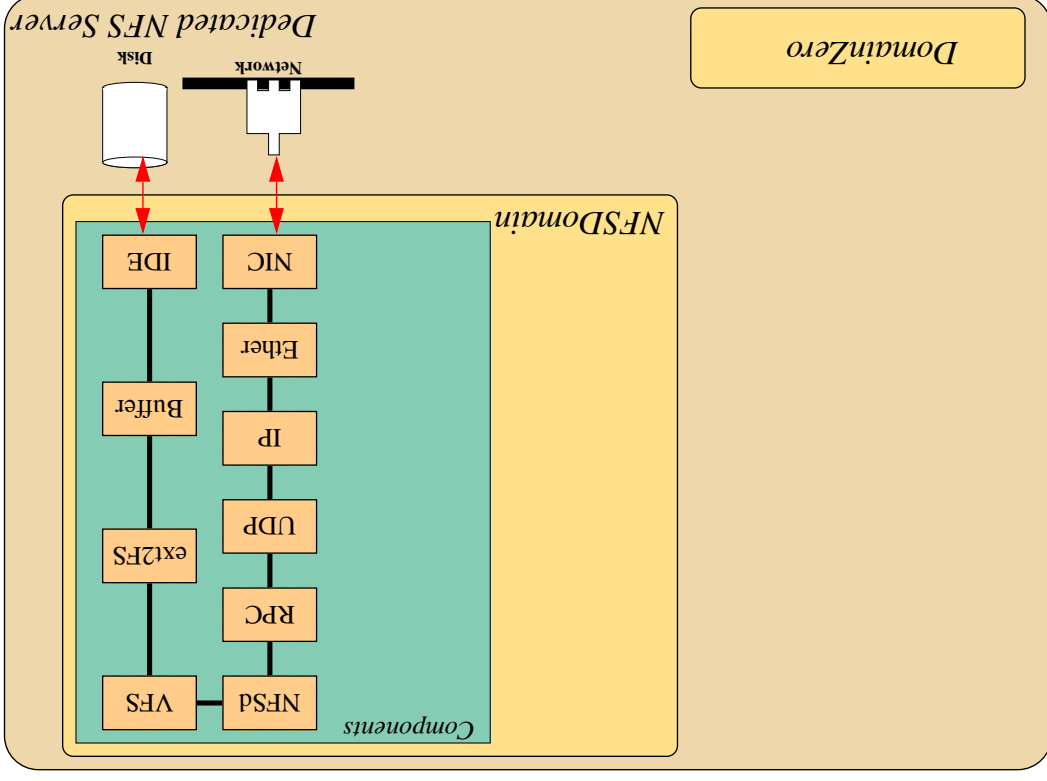
- Interrupts on the interrupted CPU are blocked during execution of the interrupt handler
- Verifier checks interrupt handler for upper limit of execution time
 - can insert runtime checks to ensure timely termination
 - runtime check can terminate interrupt handler and initiate counter measure (e. g.,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

Protection: Device Driver

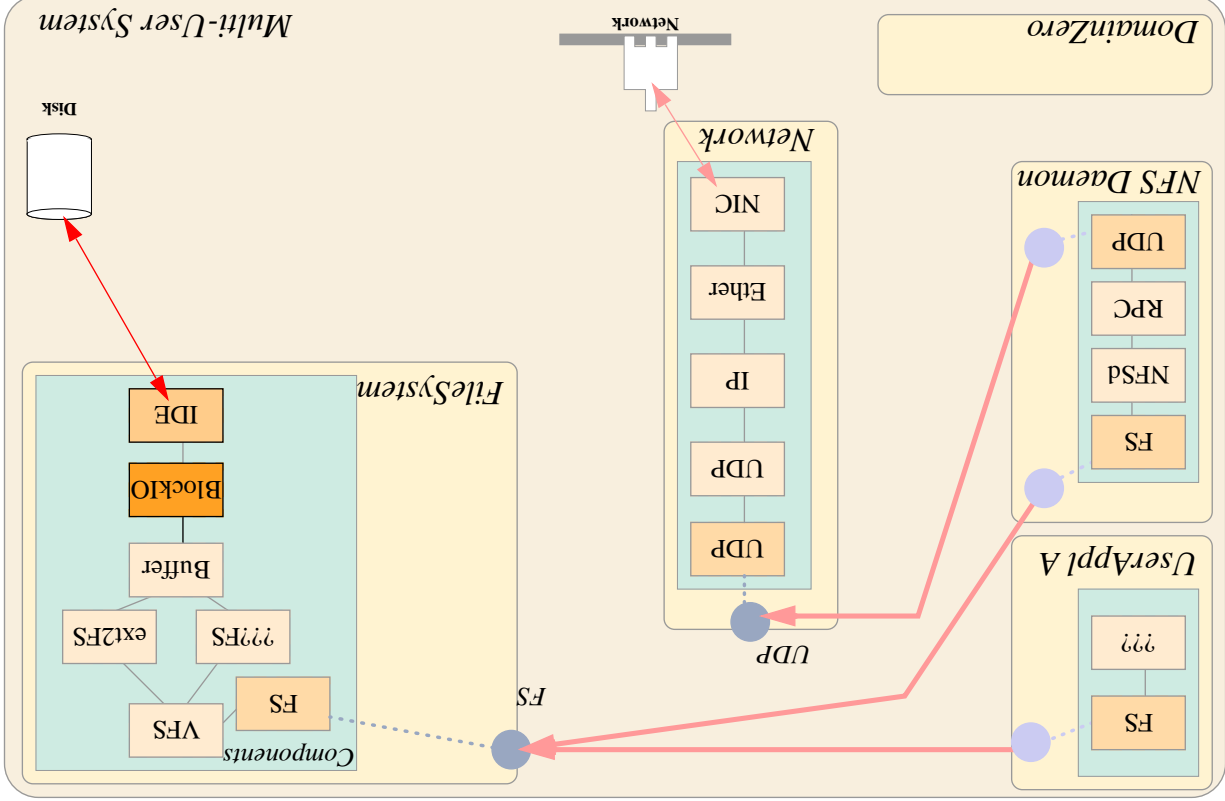
- X protection is based on type safety and portals
- Some domains can circumvent these mechanisms
 - DomainZero, Translator, Verifier → Trust
 - (some) device drivers



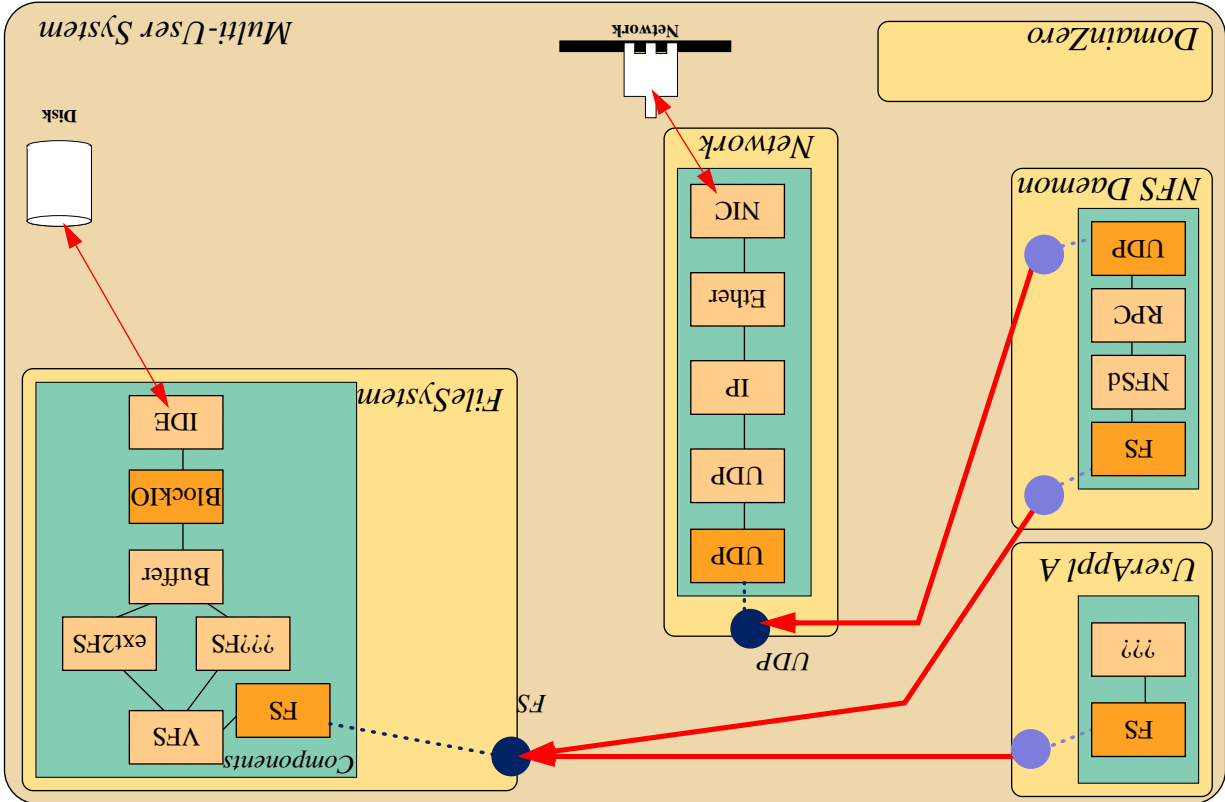
Building an OS: A Dedicated System



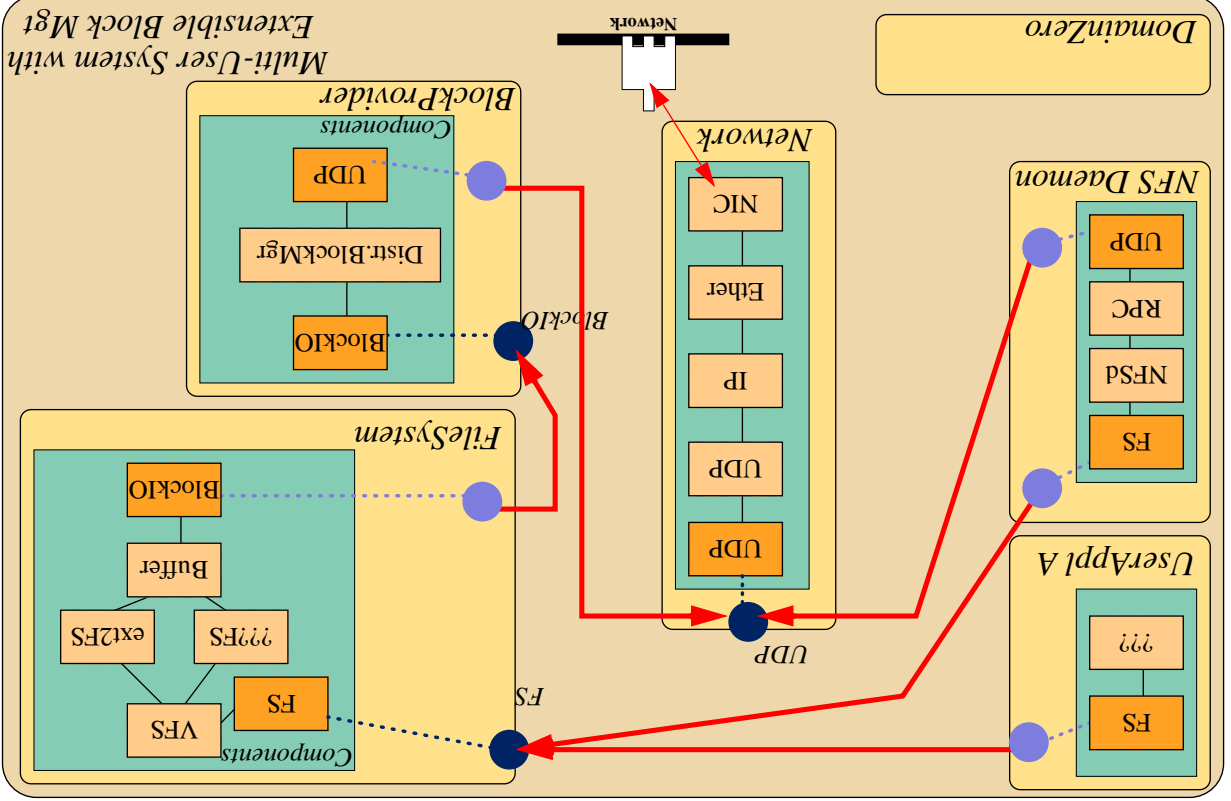
Building an OS: A Multiuser System



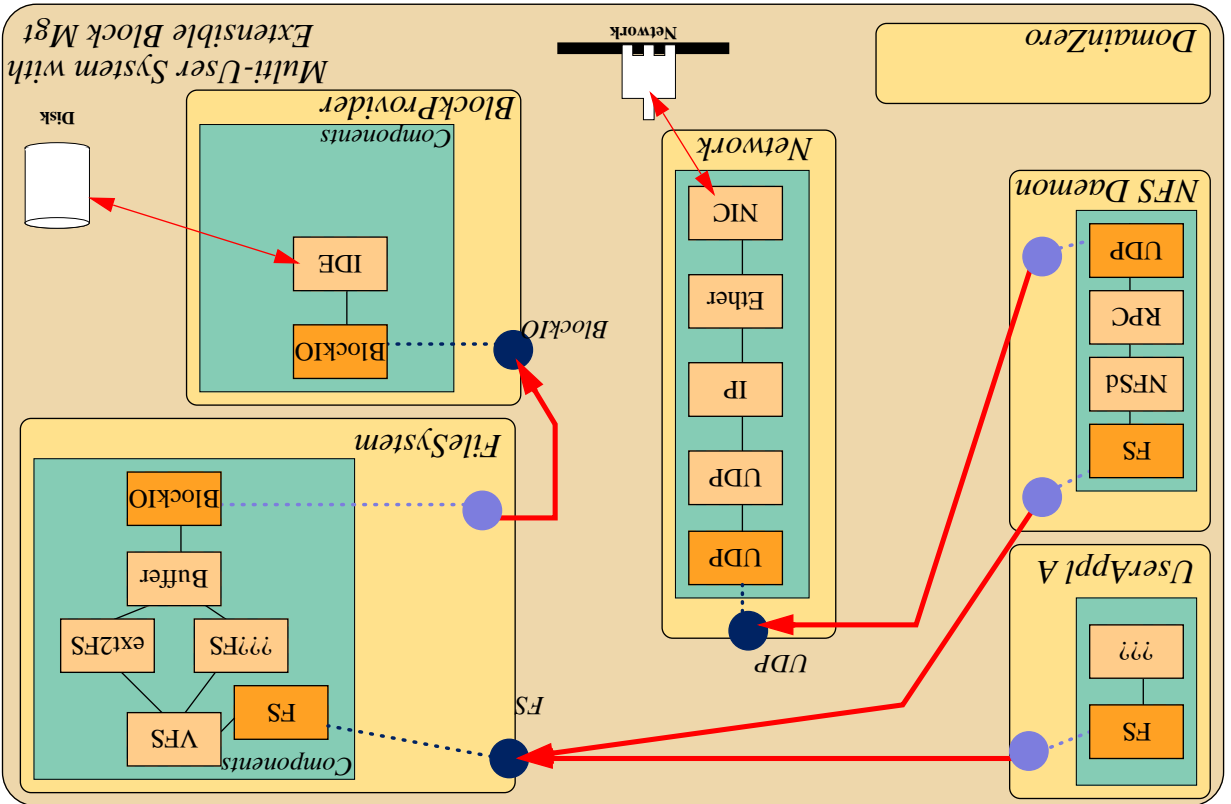
Building an OS: A Multiuser System



Building an OS: Extensibility



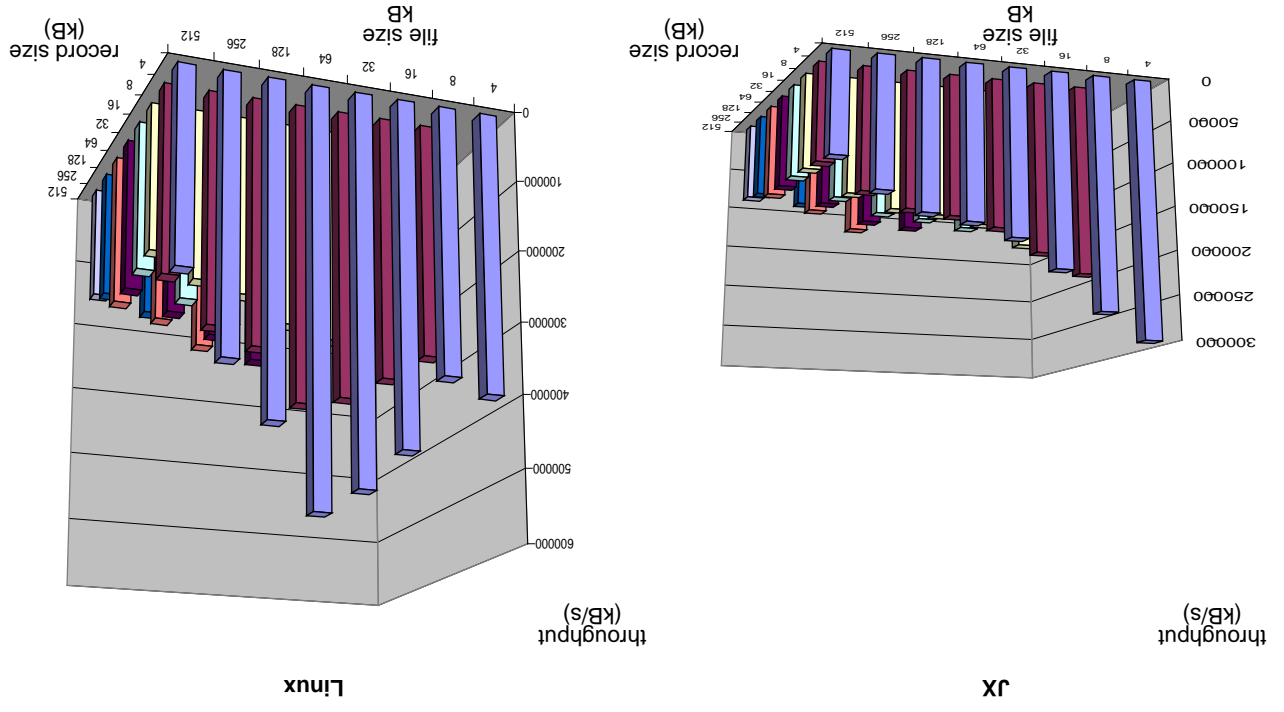
Building an OS: Extensibility



Performance

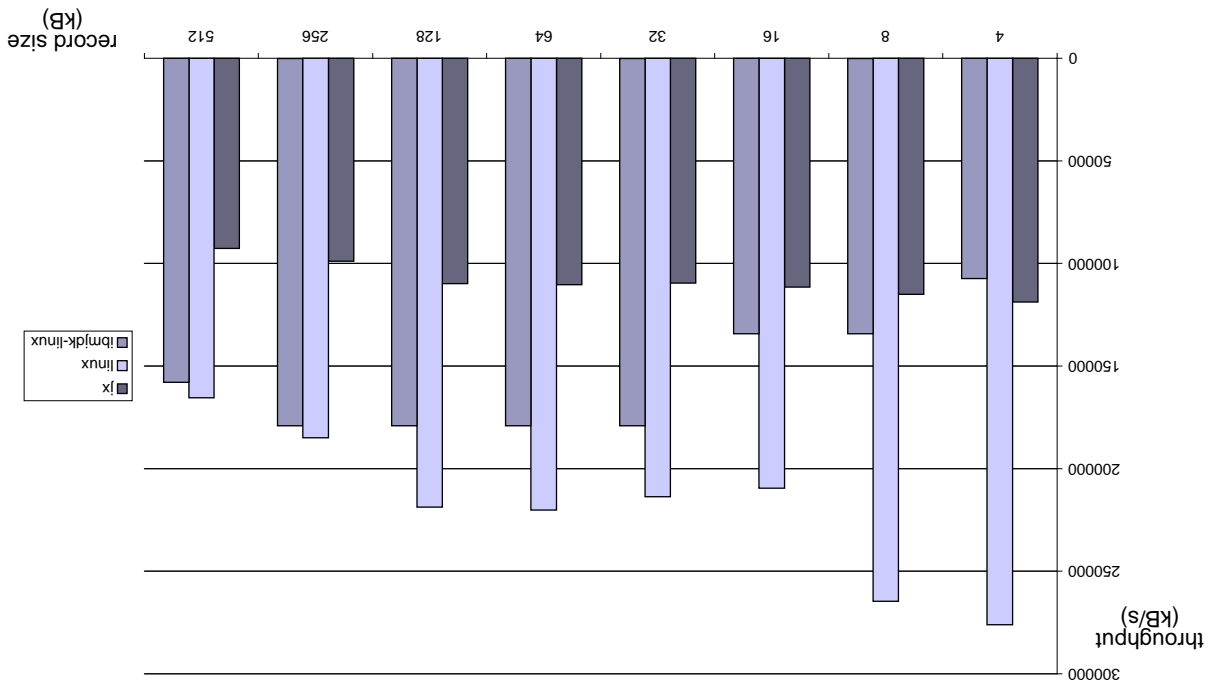
- Hardware:
500MHz PIII, 128MB RAM,
IDE: Maxtor 91303D6, 12427MB, 512KB Cache
NIC: 3C905B 100 MB/s
- IPC:
 - Portal round trip 650 cycles
 - L4 (including RPC stubs): 800 cycles
 - KaffeOS: 27270 cycles

■ Iozone-like benchmark: re-read



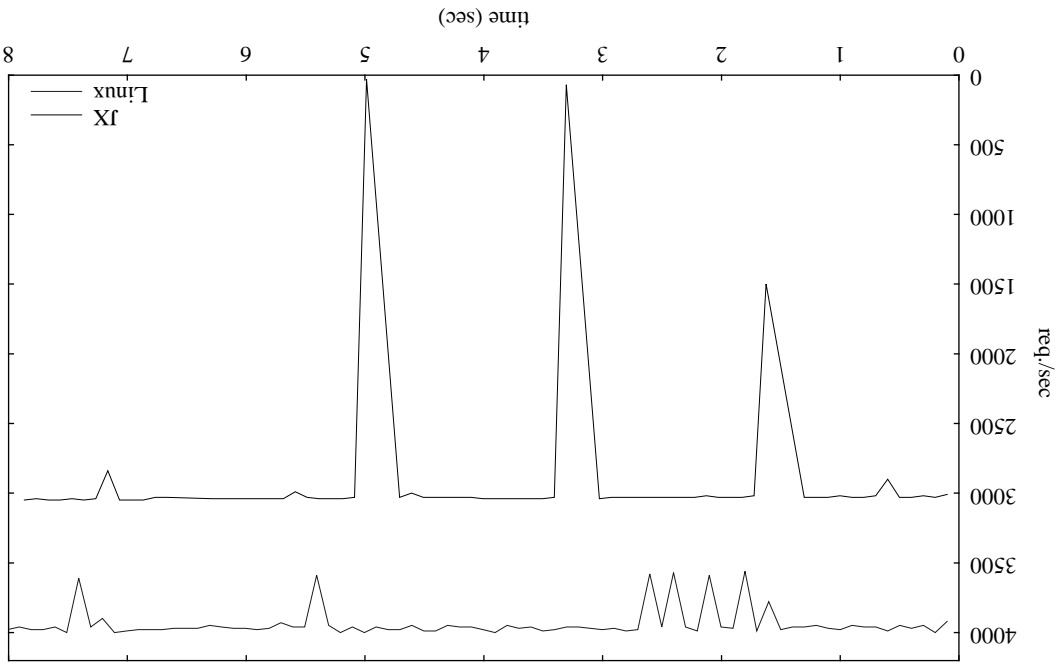
Performance

iozone-like benchmark: re-read of a 512 kB file



Performance

JX as NFS server: getattr request rate



Performance: JX Advantages and Limitations

■ Advantages:

- No expensive border crossings (JNI, OS border)
- Safe inlining of OS-level code into application code
- Avoid locking in favour of specialized scheduling

■ Limitations:

- Unavoidable safety checks
- Semantic gap between stack machine and register machine

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

Component Interaction

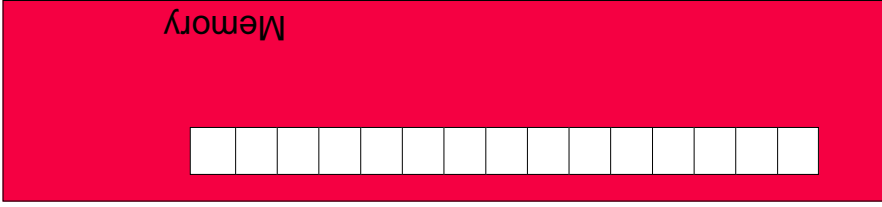
Flexibility

co-located components	co-located components	dis-located components
parameter passing	by reference possible	parameter objects must be copied
thread	can execute in same thread	threads must be switched
resources	caller can be trusted to care-fully use resources (e. g., Memory objects)	access rights for Memory objects must be restricted

- Scheduling: scheduling strategy of all co-located components must be compatible
- Execution engine of co-located components must be identical (translated or interpreted)

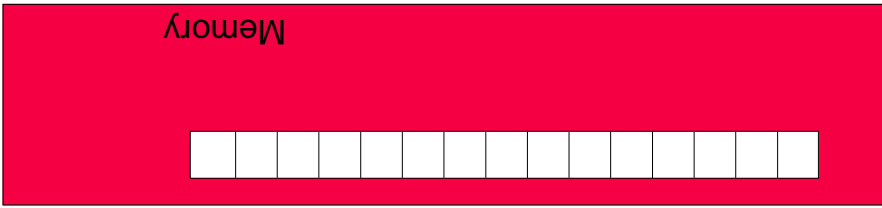
int a
byte b
int c
char d

class X



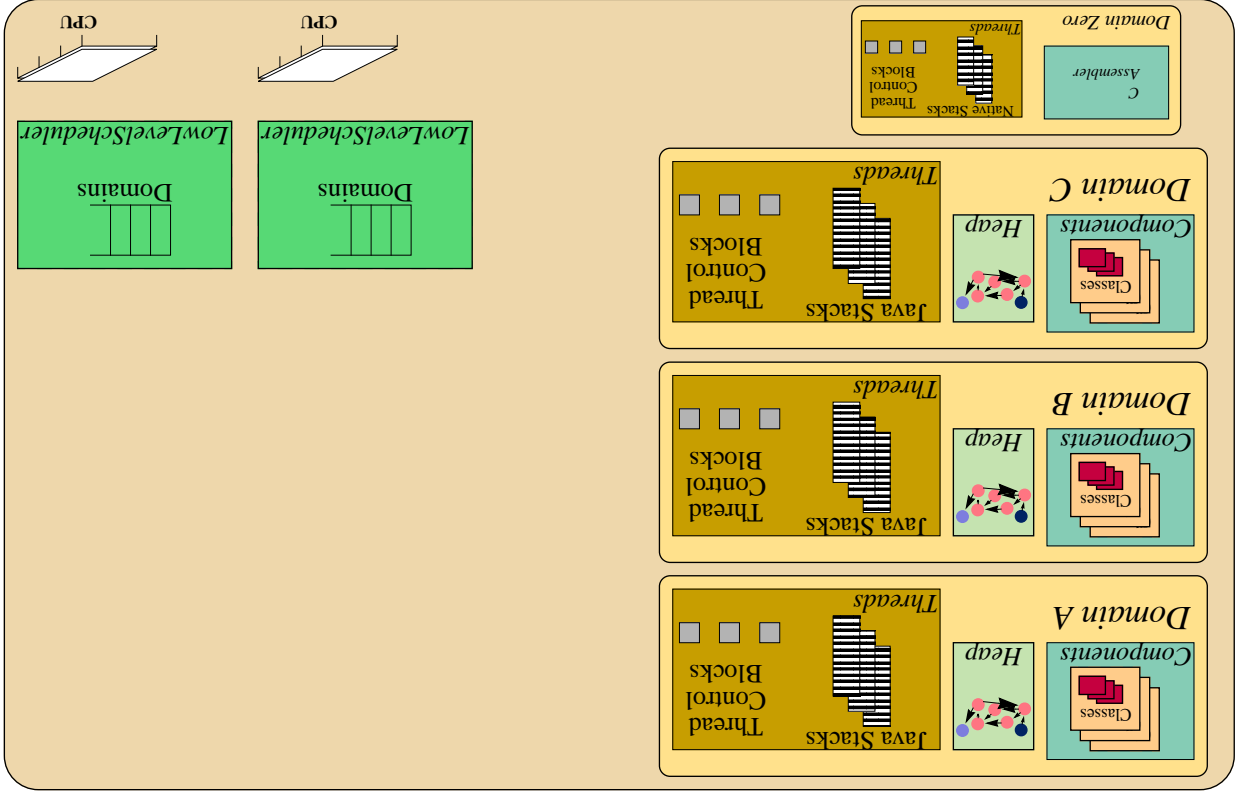
- Performance problem: range check
- (partial) solution: map objects to memory range

Memory Mapping

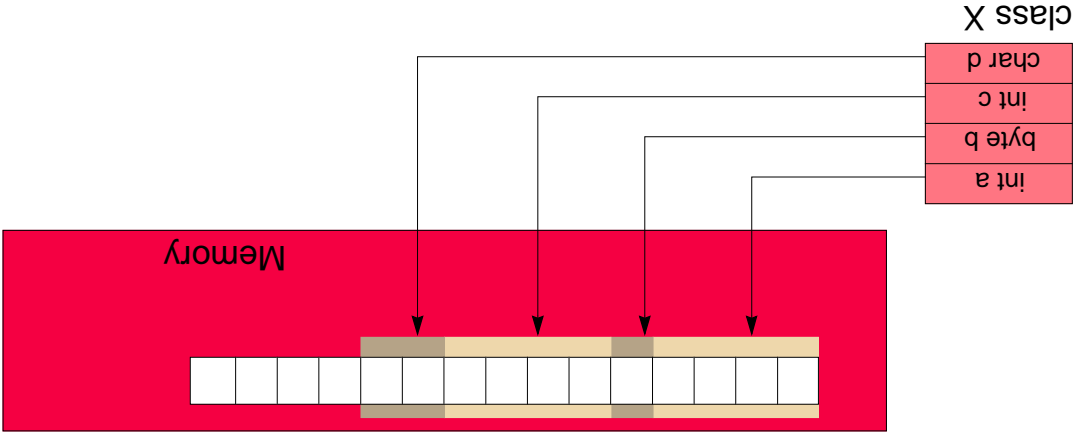


- Performance problem: range check
- (partial) solution: map objects to memory range

Memory Mapping



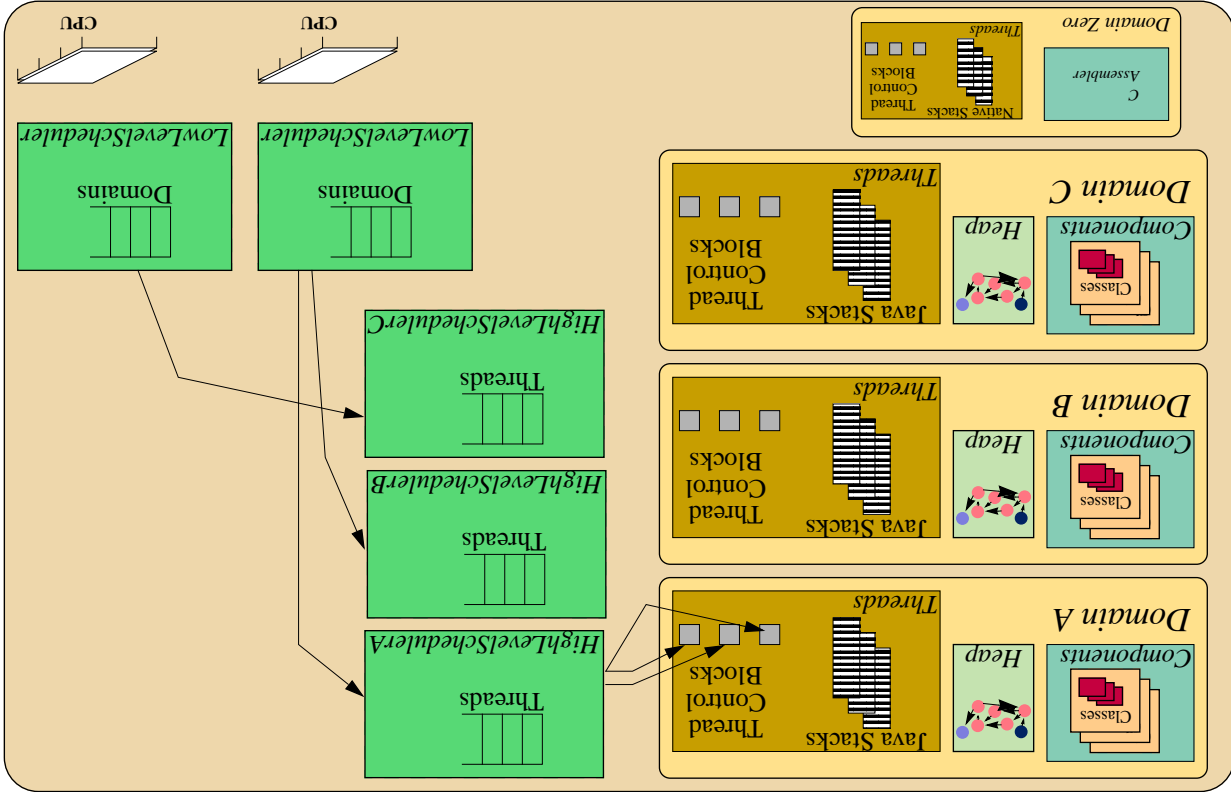
Scheduling



- Performance problem: range check
- (partial) solution: map objects to memory range

Memory Mapping

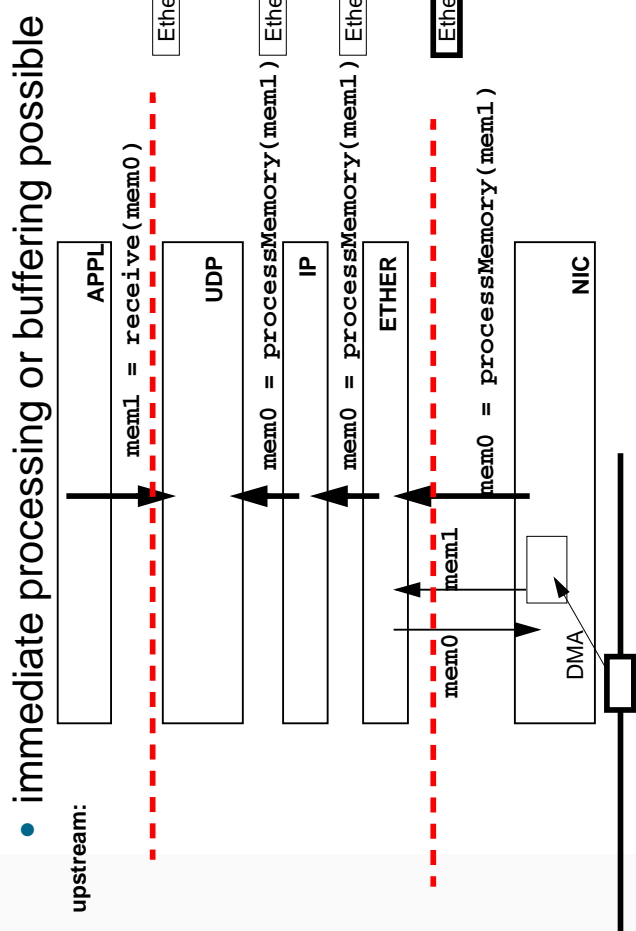
Scheduling



Resource Management

- Physical resources (CPU, Memory)
 - resource management is supported by DomainZero (but no policy)
- Device-specific physical resources (e. g. network bandwidth)
 - completely managed by a domain
- Virtual resources
 - managed by their respective domains (e. g. TCP port numbers, files)

Zero-copy using Memory



© 2001 Michael Golm, Universität Erlangen, Informatik 4

Hardware Access

