

Instructor Notes

Exokernel: An Operating System Architecture for Application-level Resource Management

Dawson R. Engler, M. Frans Kaashoek, and James O'Toole Jr.

In the Proceedings of the 15th ACM Symposium on Operating Systems Principles (SOSP '95), Copper Mountain Resort, Colorado, December 1995, pages 251-266.

1. What is the overriding goal of exokernel? How is this different than what previous extensible systems did (Nucleus, Mach, HYDRA)? How is this different from a VVM?

- Separate protection from management
- lowest level primitives possible:
 - more efficient
 - more flexibility

export
hardware
resources

- Similar goals, but more extreme

Compare - don't try to abstract resources

or virtualize them (like VMm)

- not even "mechanism"

2. What are exokernel's design principles?

- securely expose hw
 - instr.
 - resources
- (avoid resource management)
let libOS do
- expose allocation
 - let apps request physical resources
(don't make implicit)
 - this CPU?
 - this page?
- expose names
- expose revocation
 - if something taken away, let app know

3. Where does policy belong? How does exokernel handle conflicts/competing applications?

• in libOS

• if managed by same libOS \Rightarrow fine

• otherwise??

"must contain policy to arbitrate between competing libOSes"

4. What is the purpose of a **secure binding**? Why can secure bindings achieve good performance? How can secure bindings be used to multiplex physical memory? To multiplex the network?

- guard resources efficiently
 - decouples authorization from use
- 1) Primitives for setting up should be good match w/ hw (or exokernel)
 - 2) only need to be ~~done~~ ^{setup (and understand semantics)} @ bind, not access
-
- Mem: check capability to access page on TLB miss (fast access)
 { binding
 - Network:
 multiplex network requires interpreting incoming msgs
 → download "safe" code into kernel
 to do this interpretation

5. What is an Application-Specific Safe Handler (ASH) and why is it useful?

- Code run in kernel on behalf of app
- Faster who switches out

6. Why is resource revocation visible to applications using an exokernel?

- Traditional OS: invisible - just take timeslice or page away
- Exo: believe app should be able to respond (save reg on context switch)
(~~stop~~ out w/ phys. mem)

7. Why is an **abort protocol** needed?

- If libOS doesn't revoke as desired
 - 1) Could kill libOS
 - 2) Tell libOS gone + guarantee some minimal resources

8. How did the authors demonstrate the flexibility of the exokernel architecture?

1) Aegis is efficient w/ base costs


Section 7} ExOS

- extensible RPC (who saves rep)
- inverted page table (instead of linear)
- extensible scheduler
 - stride

9. How can various resources be "multiplexed" without policy decisions?


point of mux is decision point!

2)




should they get
to share?

1)



who, when? policy!

3)



have to ~~download~~ run packetfilter
code

10. Conclusions?

Other problems?

- large executables
- portability??

What was going on in OS world?

- OS was viewed as couldn't be changed, very complex
- HP-UX, AIX, Solaris, Windows...
- Just let ~~non~~ interested developers change source code directly of Linux?

make