Adaptable Partitioning with a Real-Time Separation Kernel

Henrik Karlsson (henrik10@kth.se)

KTH Royal Institute of Technology

May 22, 2025

What is a Kernel?

Core of the Operating System

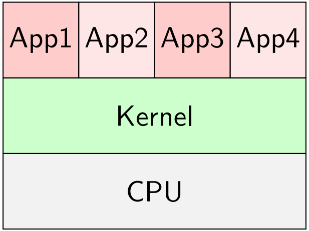

App1 App2 App3 App4

Kernel

CPU

Manages resources and provides services to applications.

What is a Kernel?

Core of the Operating System

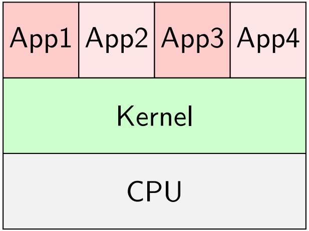

App1 App2 App3 App4

Kernel

CPU

Manages resources and provides services to applications.

“The Janitor of the CPU”

Traditional Kernels: Challenges

Windows and Linux Kernels: General-Purpose Design

Traditional Kernels: Challenges

Windows and Linux Kernels: General-Purpose Design

These kernels act as a **“multitasking janitor”**, prioritizing throughput.

Traditional Kernels: Challenges

Windows and Linux Kernels: General-Purpose Design

These kernels act as a **“multitasking janitor”**, prioritizing throughput.

- **Security vulnerabilities** – Doors left unlocked.
- **Performance bottlenecks** – The janitor is overwhelmed.
- **Safety issues** – A worker monopolizes resources.
- **Information leakage** – Sensitive data is not erased.

Our Solution: Capability-based Partitioning Kernel

S3K: A Dynamic Partitioning Kernel

- Partitions the system into secure compartments.
- Uses **capabilities** for dynamic compartmentalization.

Our Solution: Capability-based Partitioning Kernel

S3K: A Dynamic Partitioning Kernel

- Partitions the system into secure compartments.
- Uses **capabilities** for dynamic compartmentalization.

~~**“The Janitor of the CPU”**~~

“The Security Guard of the CPU”

Checks tickets before granting access to resources.

Our Solution: Capability-based Partitioning Kernel

S3K: A Dynamic Partitioning Kernel

- Partitions the system into secure compartments.
- Uses **capabilities** for dynamic compartmentalization.

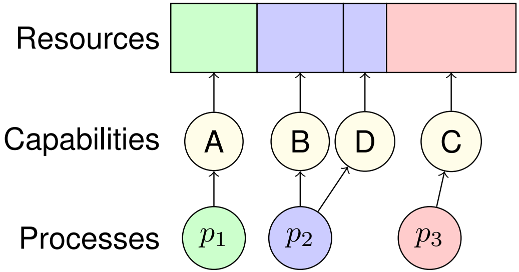

S3K Scheduler

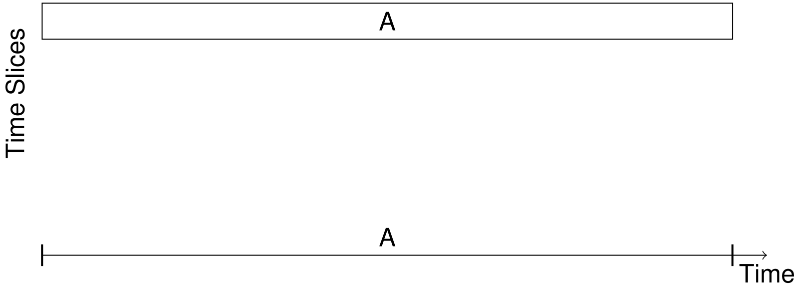

Time Slice Capabilities:

- Allocate CPU time for processes.
- CPU state is cleaned after each time slice.

S3K Scheduler

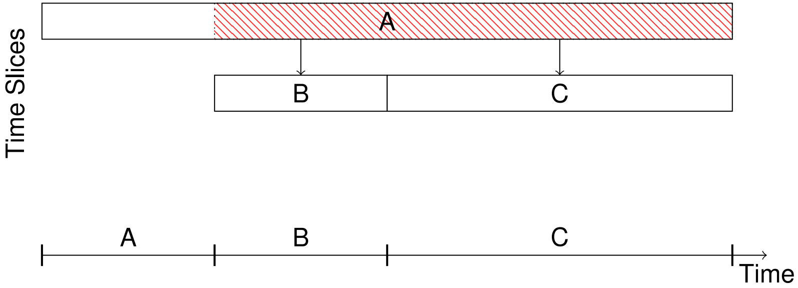

Time Slice Capabilities:

- Allocate CPU time for processes.
- CPU state is cleaned after each time slice.

S3K Scheduler

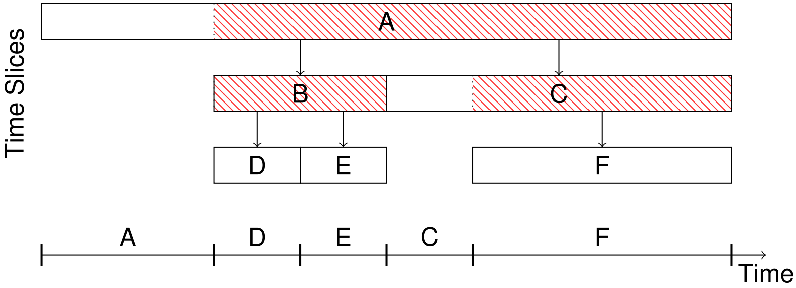

Time Slice Capabilities:

- Allocate CPU time for processes.
- CPU state is cleaned after each time slice.

S3K Capability Implementation

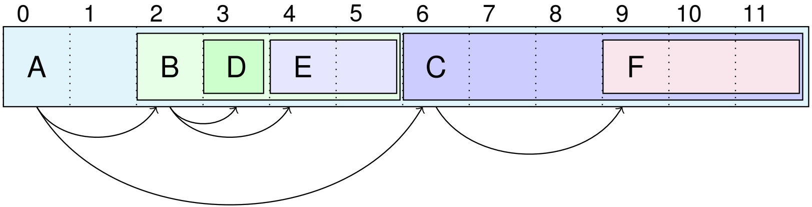

Capability Derivation Tree:

- Capability access reveal only the resources they control
- Number of child capabilities are bounded

Conclusion

S3K: Enhancing Security and Safety

- Capability-based partitioning for dynamic, secure resource management.
- Deterministic scheduling and domain management ensure safety and flexibility.
- New capability system prevents information leakage and improves system performance.