



The Formally Verified seL4 Microkernel

High-Assurance Foundation for MCS

Gernot Heiser | gernot.heiser@data61.csiro.au | @GernotHeiser

- RTCSA Keynote, Aug'20

<https://trustworthy.systems>



What Is Needed For Mixed-Criticality?

During a review process, ca Aug'17:

- [Gernot:] Temporal isolation is *necessary* for mixed criticality systems.
- [Reviewer:] Wrong, temporal isolation is *sufficient*.

What Is a Mixed-Criticality System?

“A mixed-critical system [...] supports the execution of safety-critical, mission-critical, and non-critical software within a single, secure compute platform.”
[Barhorst’09]


Criticality of a component is defined by the impact of failure:

- loss of life
- injury
- inconvenience

Certification of critical component must not depend on behaviour of less critical components
⇒ **must prevent *any* interference by less critical components!**



Preventing Interference – The OS's Job



We need an OS that can guarantee the absence of interference!

seL4: Provable Isolation



What is seL4?



The world's **first** operating-system kernel with **provable** security enforcement

World's most advanced mixed-criticality OS

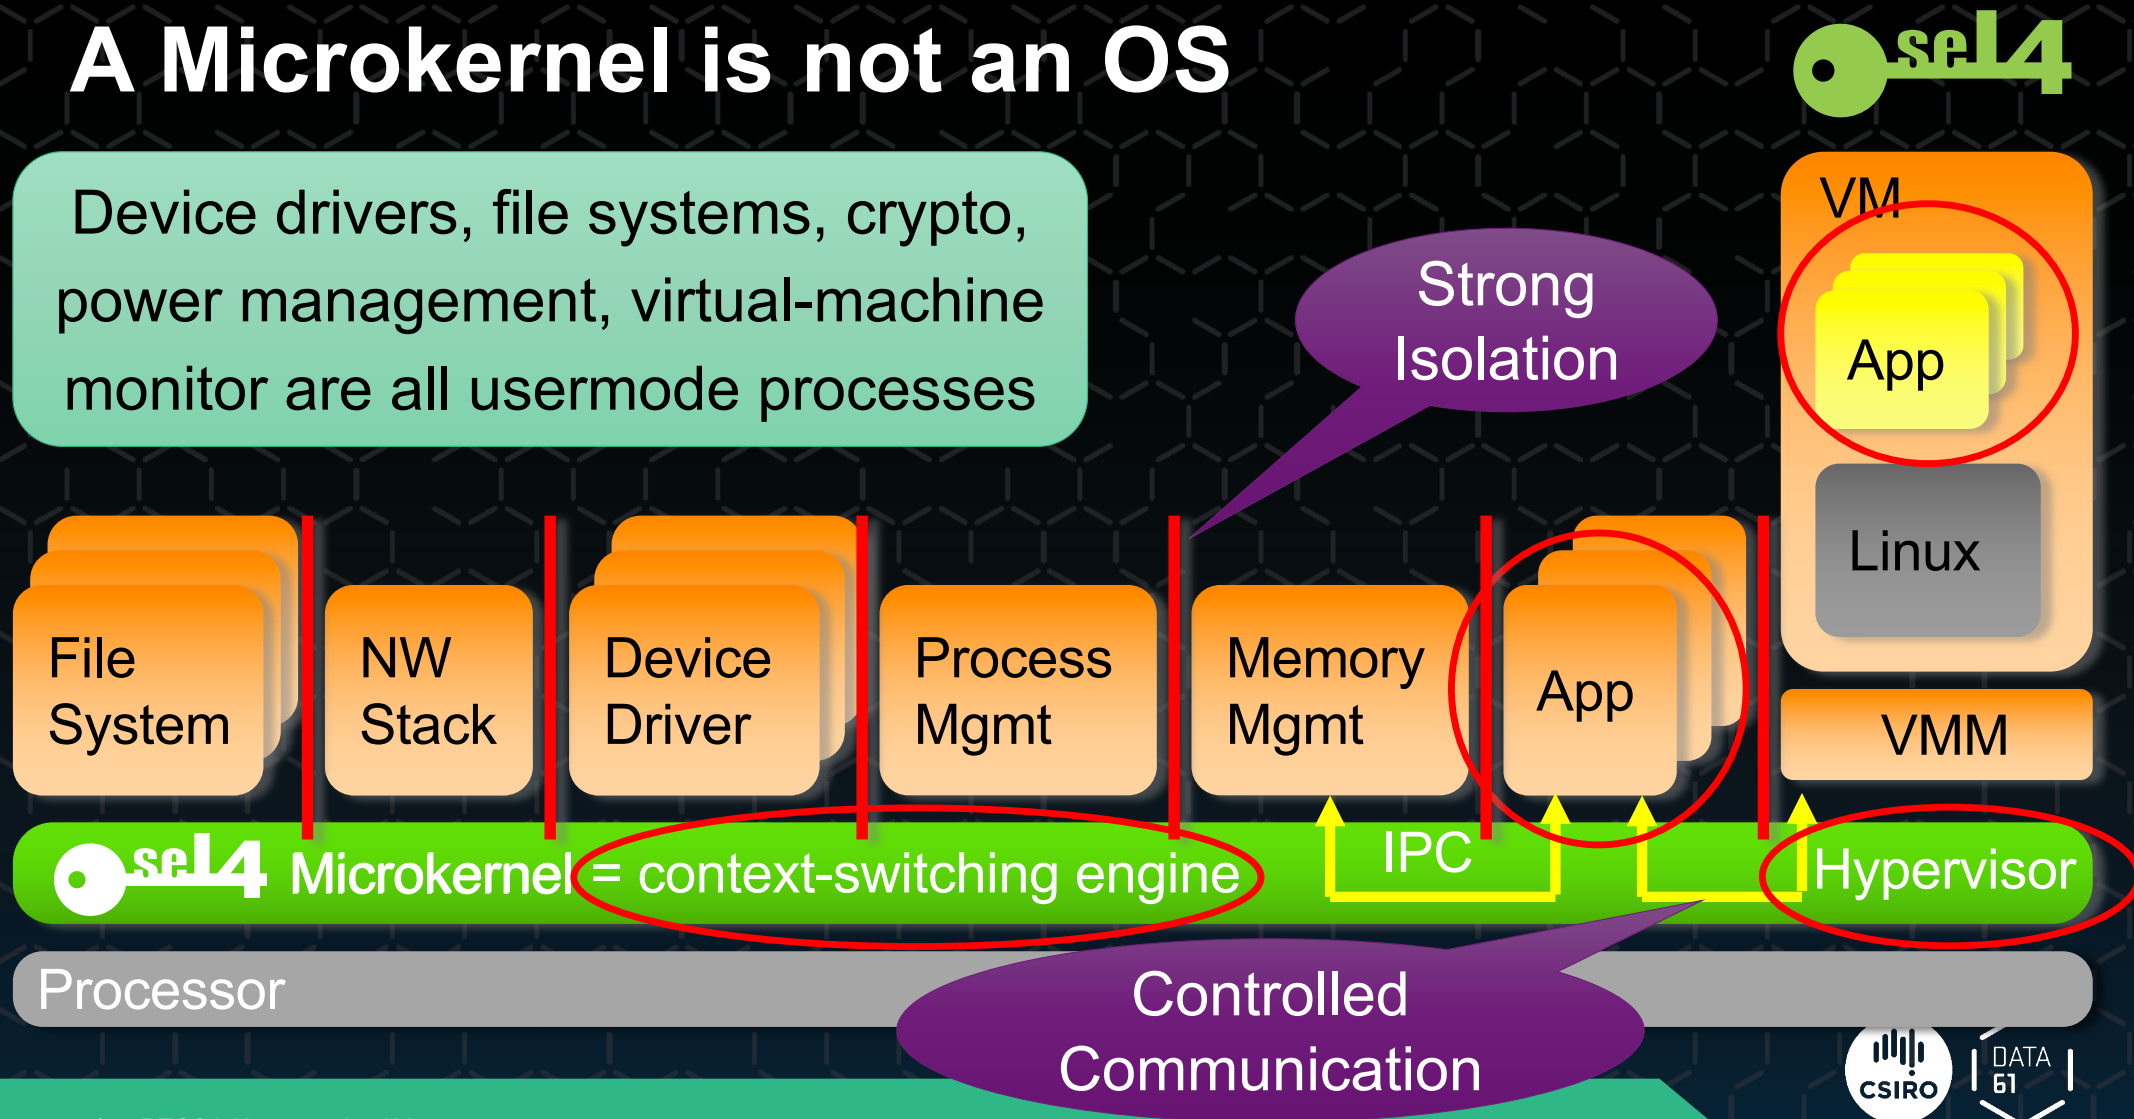
Open Source

The world's **only** protected-mode OS with complete, sound timeliness analysis

The world's **fastest** general-purpose microkernel, designed for **real-world** use

A Microkernel is not an OS

Device drivers, file systems, crypto, power management, virtual-machine monitor are all usermode processes



Capability-Based Access Control



Capability = Access Token:
Prima-facie evidence of privilege



Obj reference
Access rights

Eg. read,
write, send,
execute...



Object

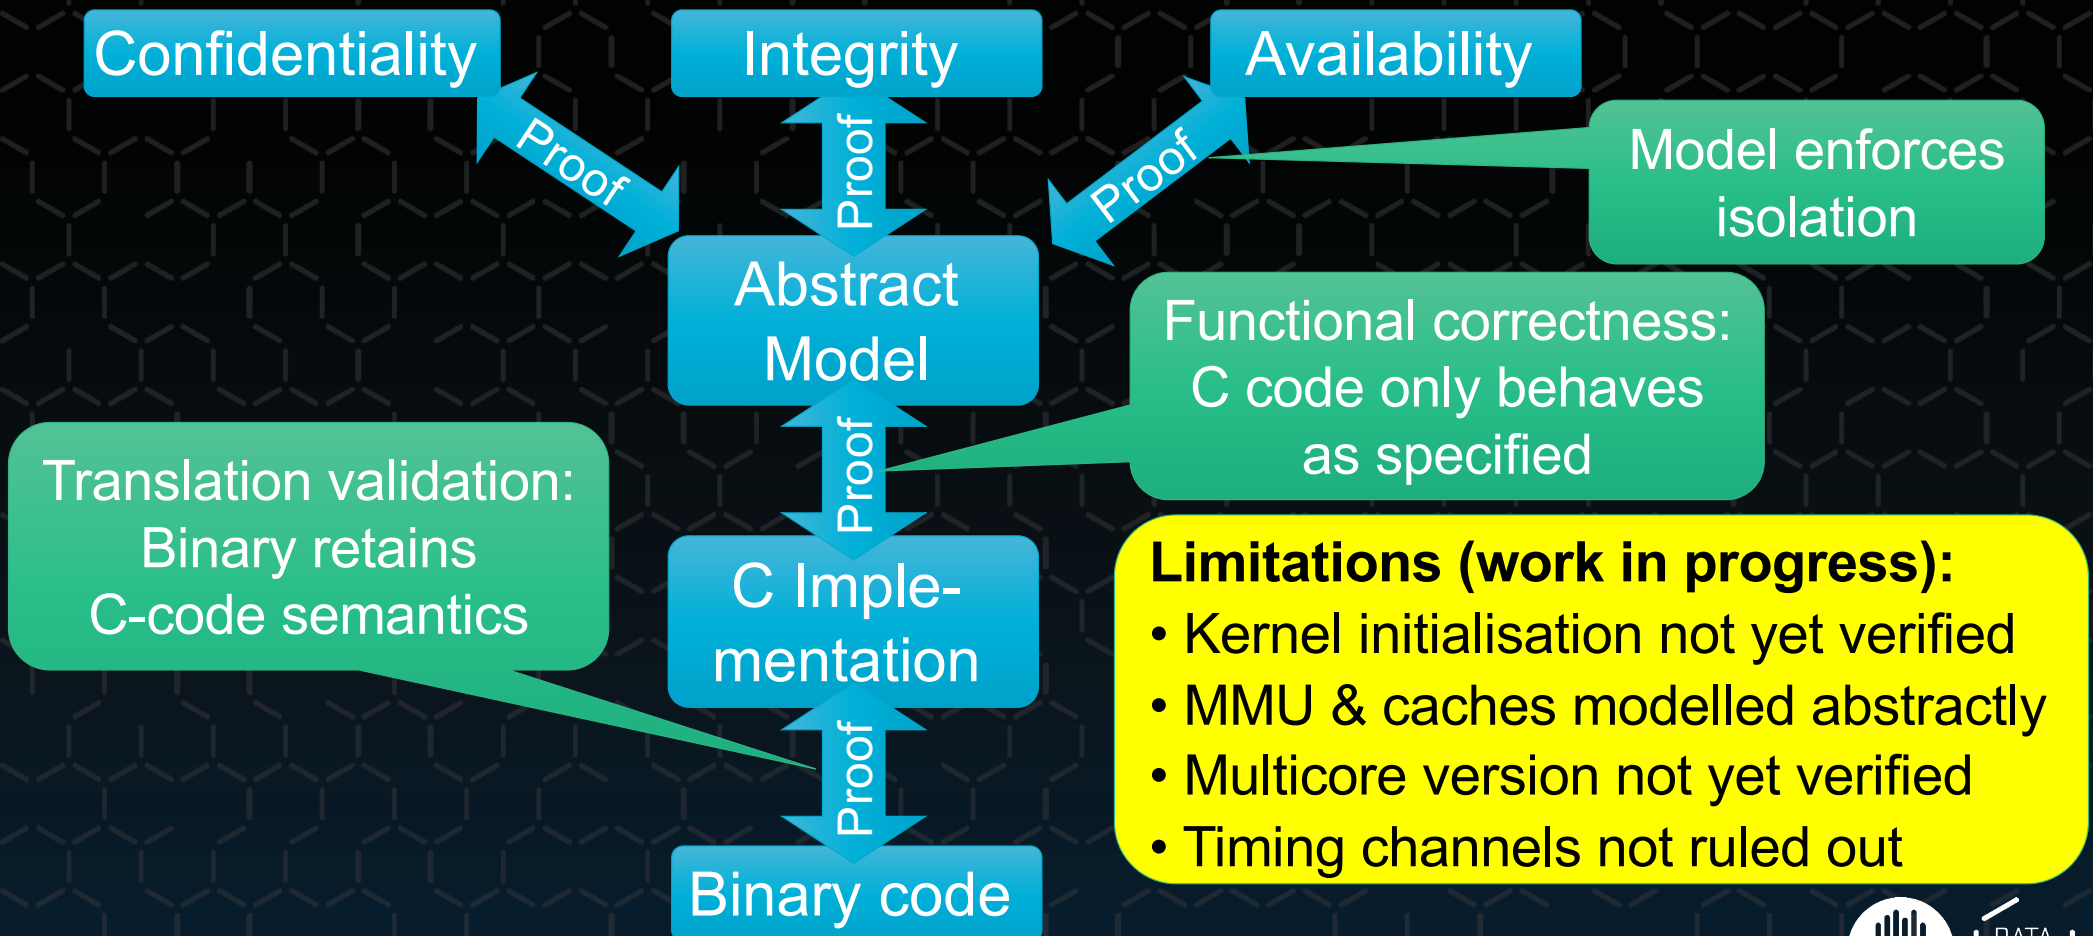
Eg. thread,
address
space

Any system call is invoking a capability:
`err = method(cap, args);`

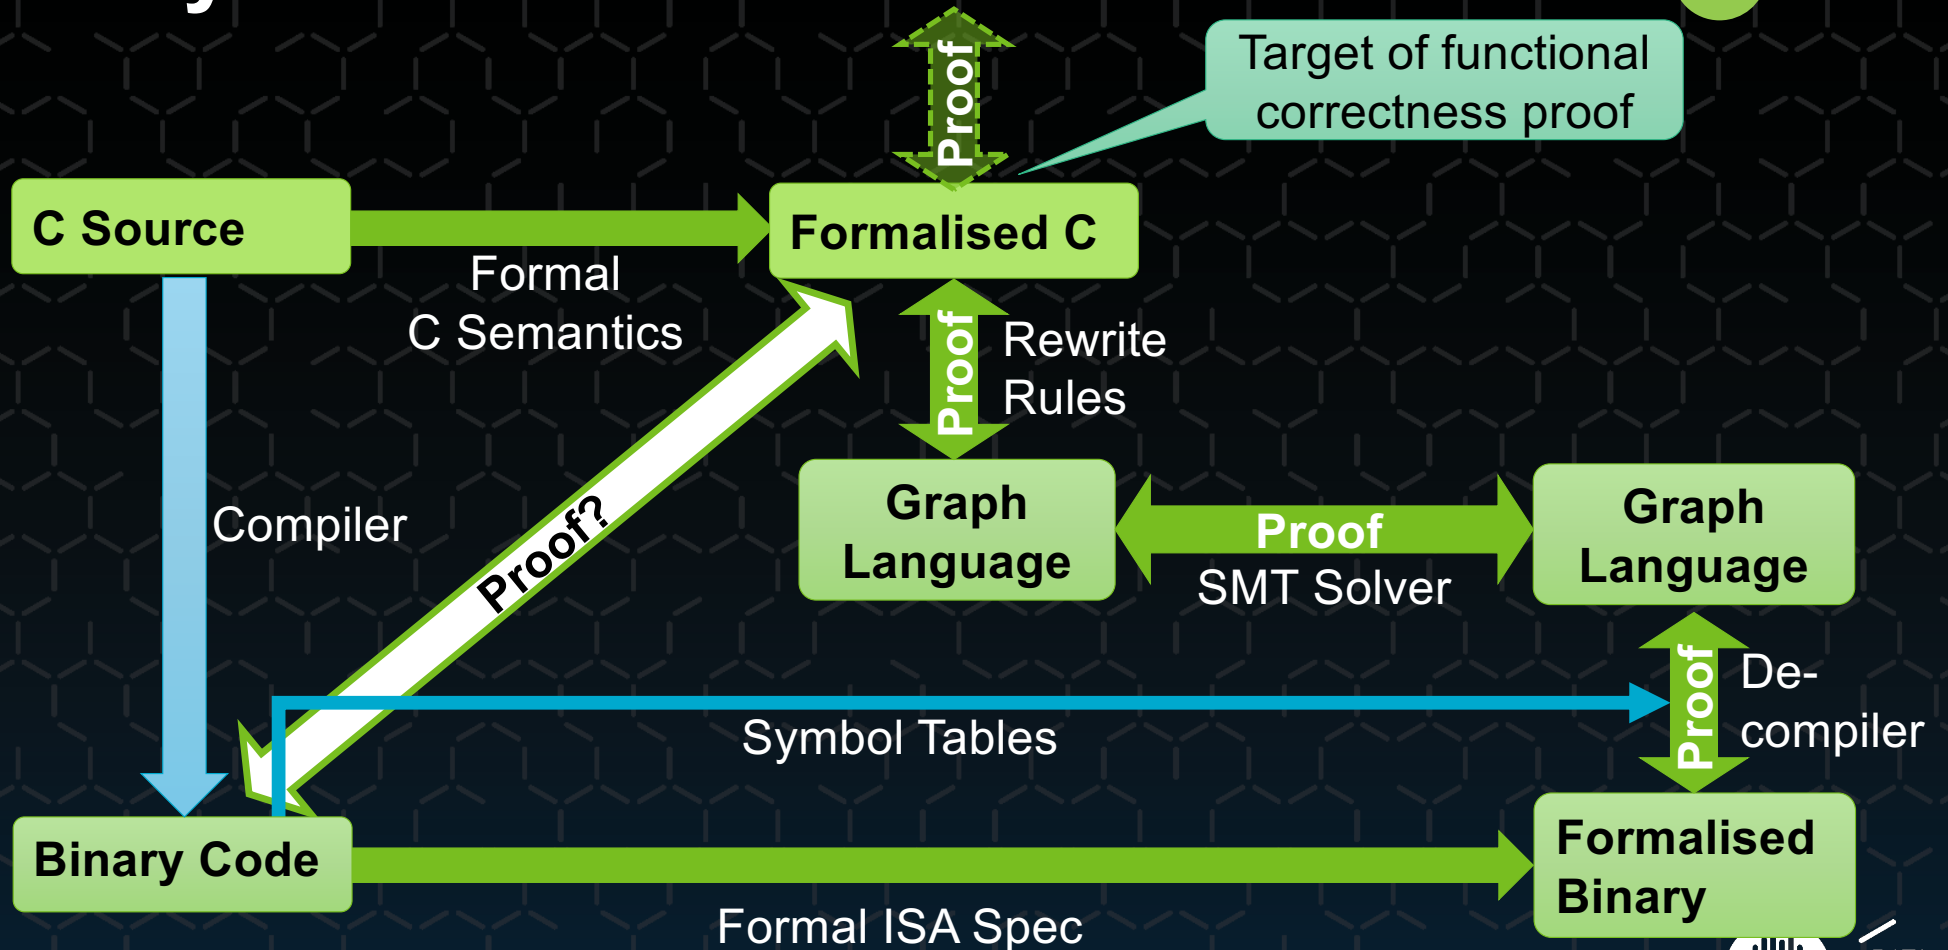
Capabilities provide:

- Fine-grained access control
- Reasoning about information flow

Proved Spatial Isolation



Binary Code Verification




Isolation by Architecture



se14




Issue: Capabilities are Low-Level



>50 capabilities
for trivial program!

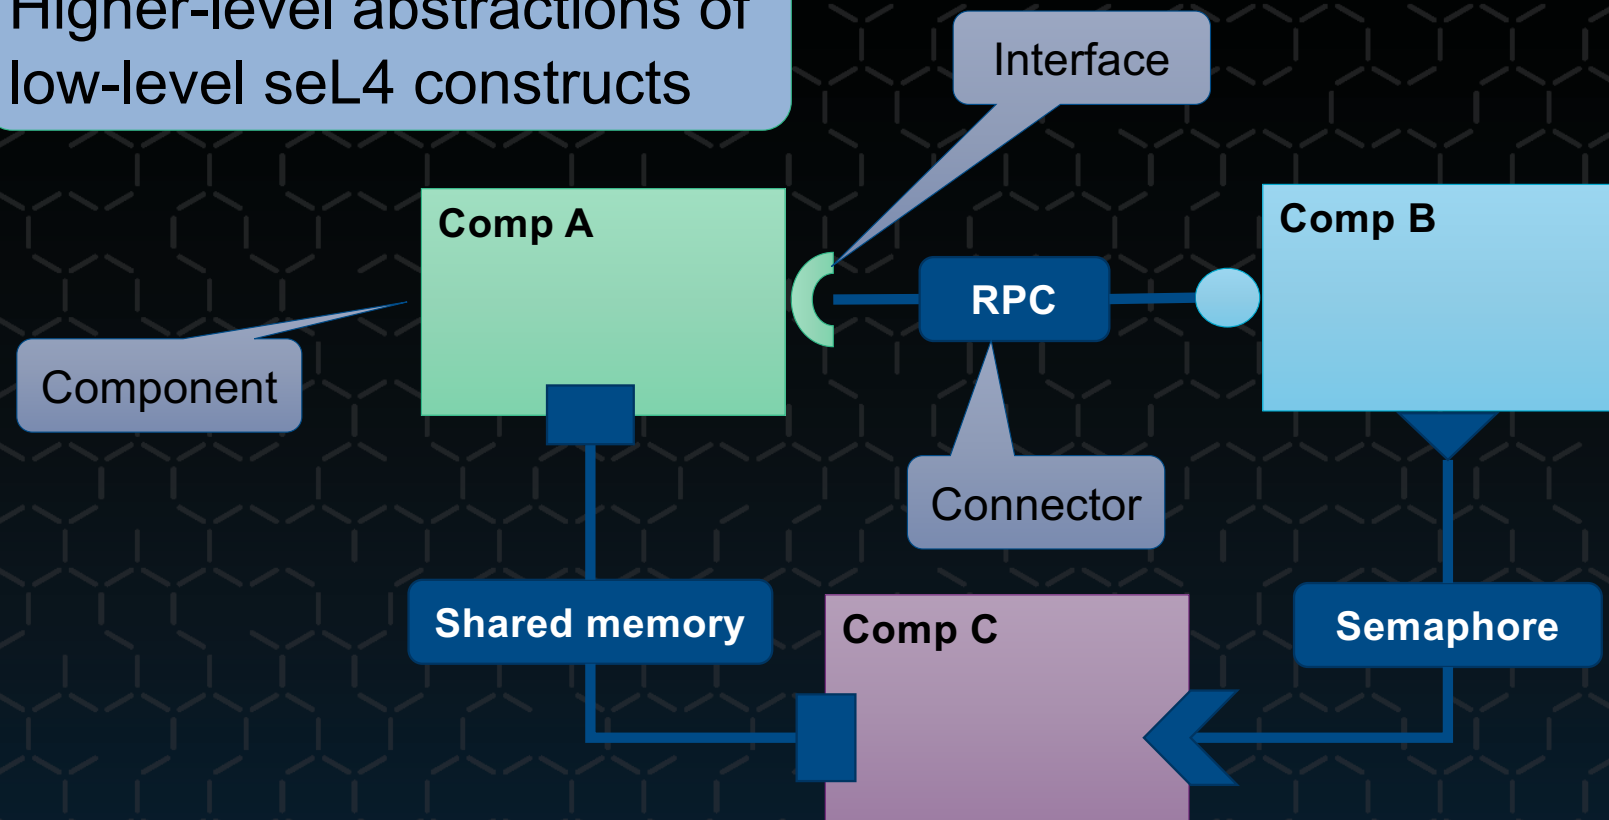
Simple But Non-Trivial System




Component Middleware: CAmkES



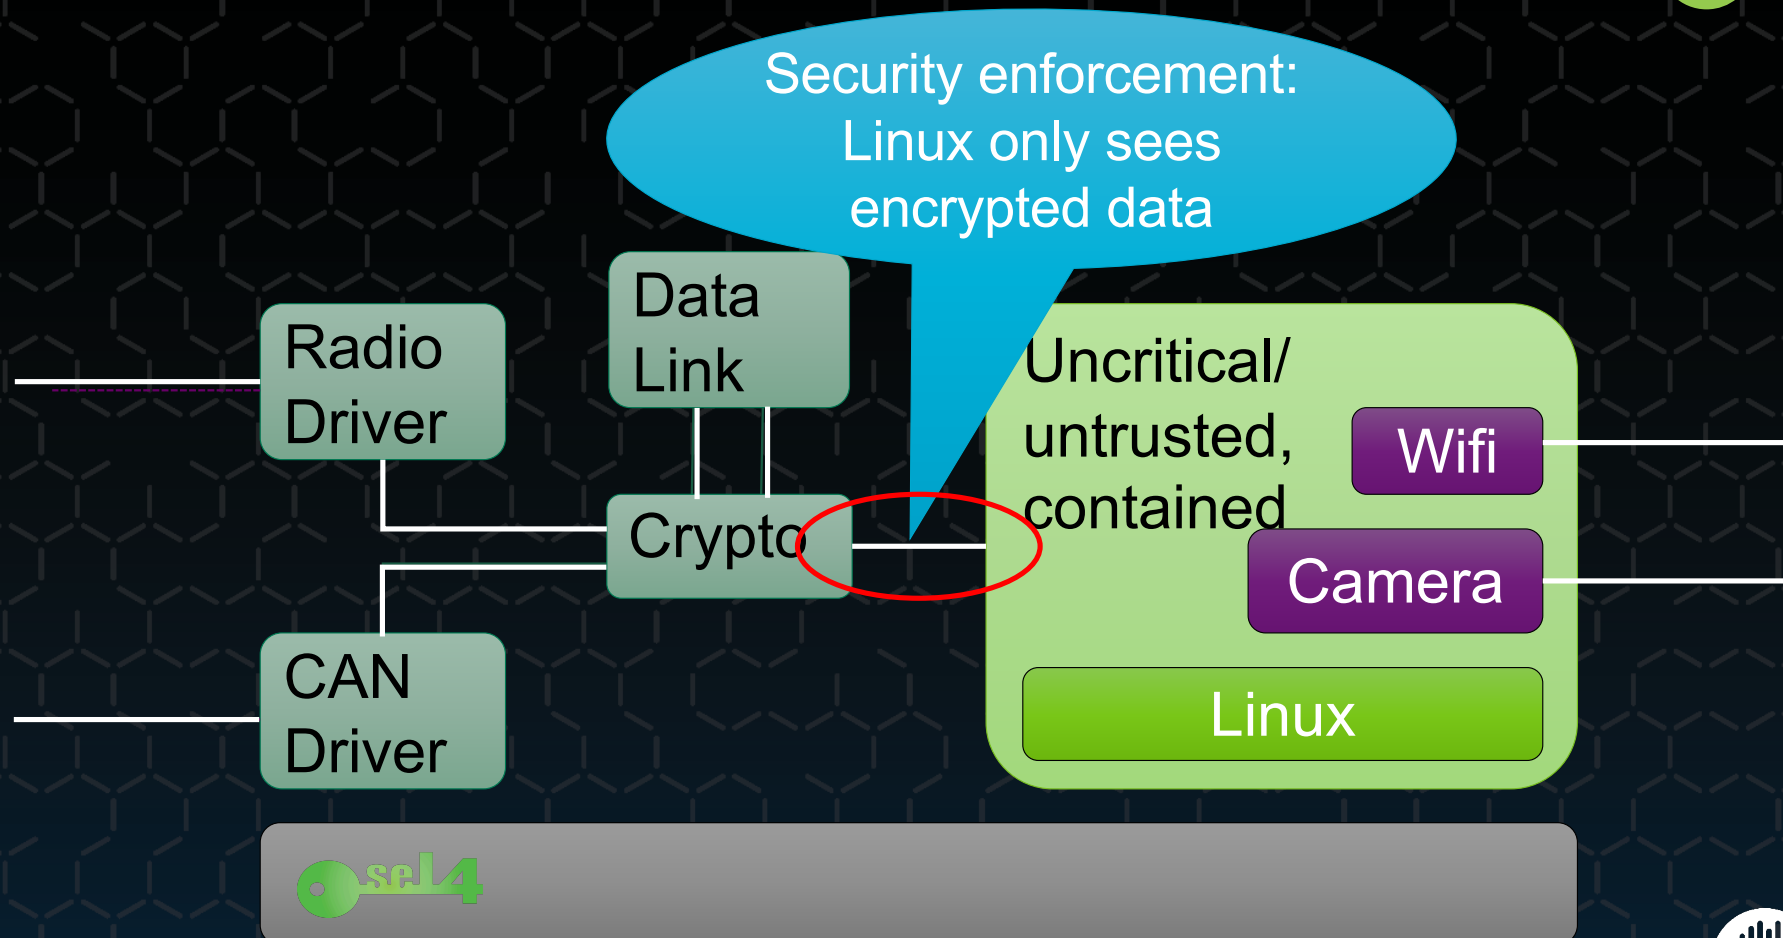
Higher-level abstractions of
low-level seL4 constructs



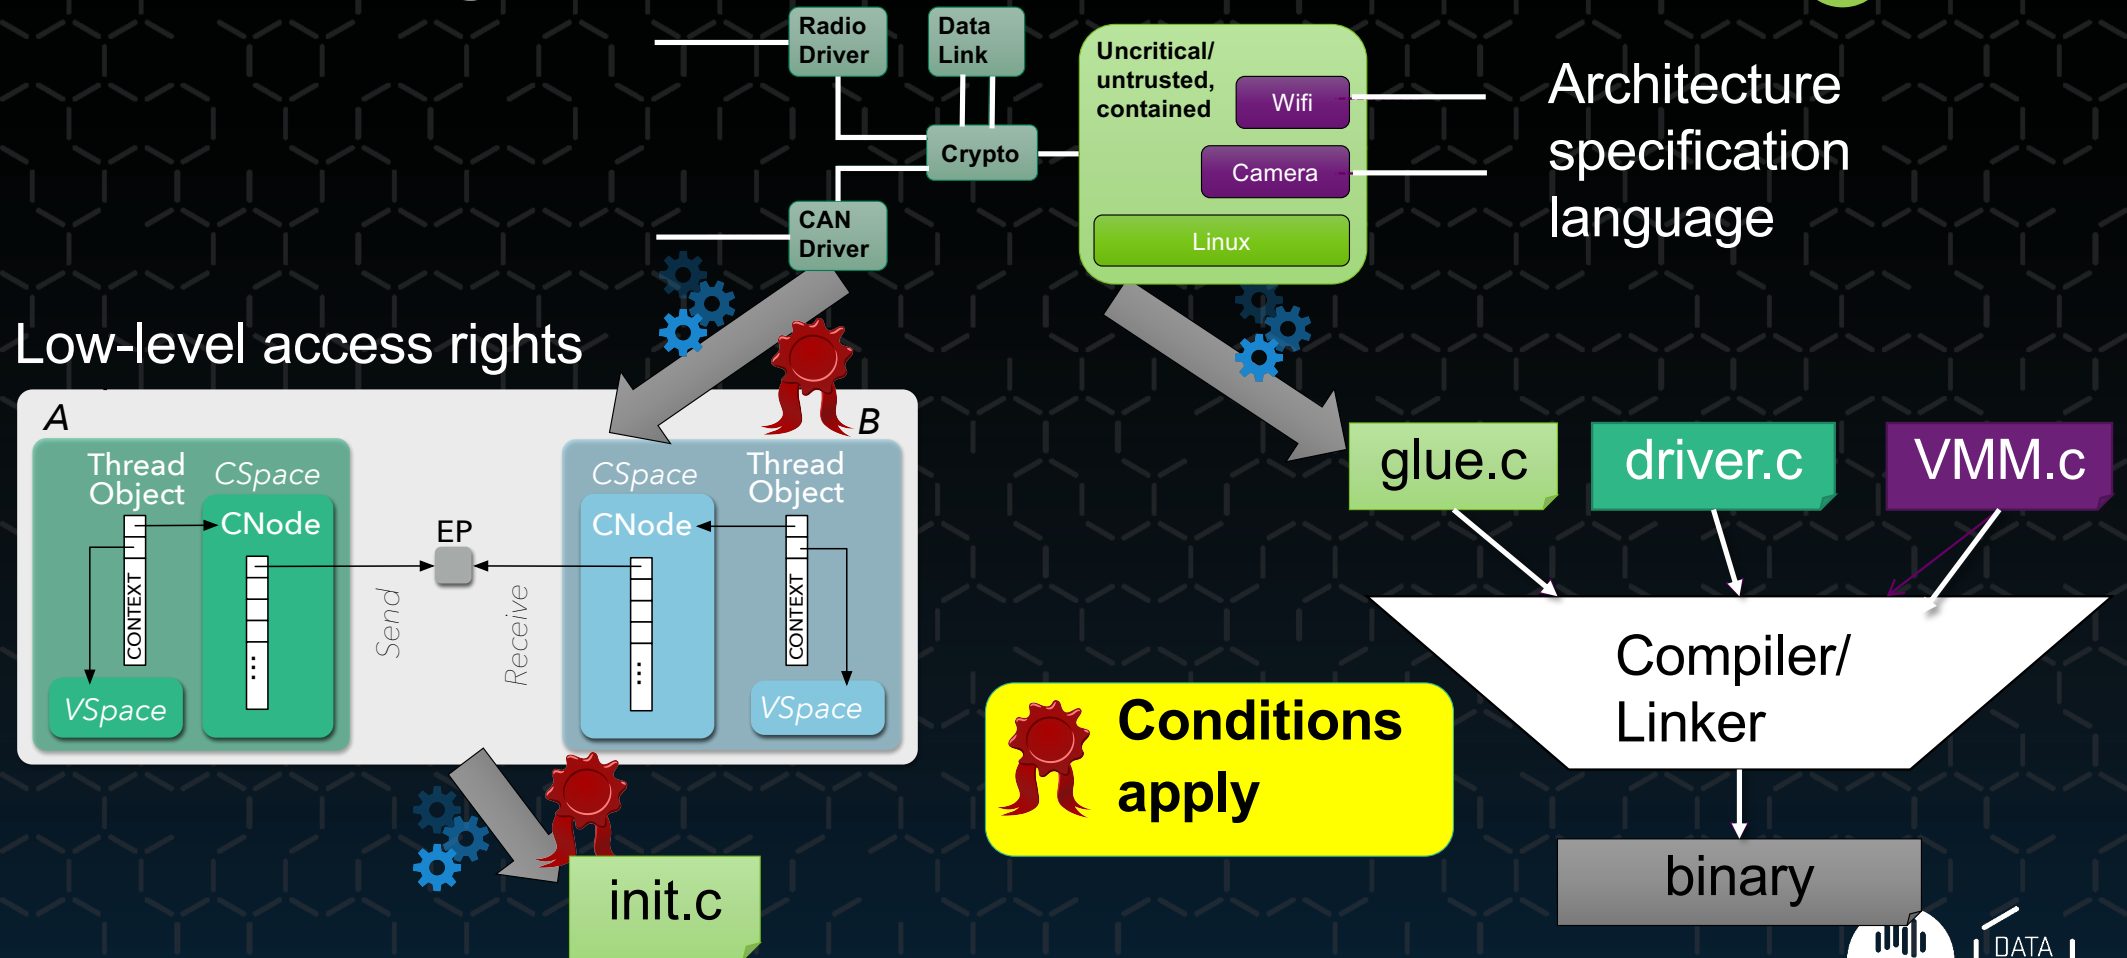
Trivial System in CAmkES



HACMS UAV Architecture



Enforcing the Architecture



Military-Strength Security



Unmanned Little Bird (ULB)

**DARPA HACMS:
Retrofit existing
system!**



Autonomous trucks

**Secure
Comms
Dongle**



**Cross-Domain
Desktop
Compositor**



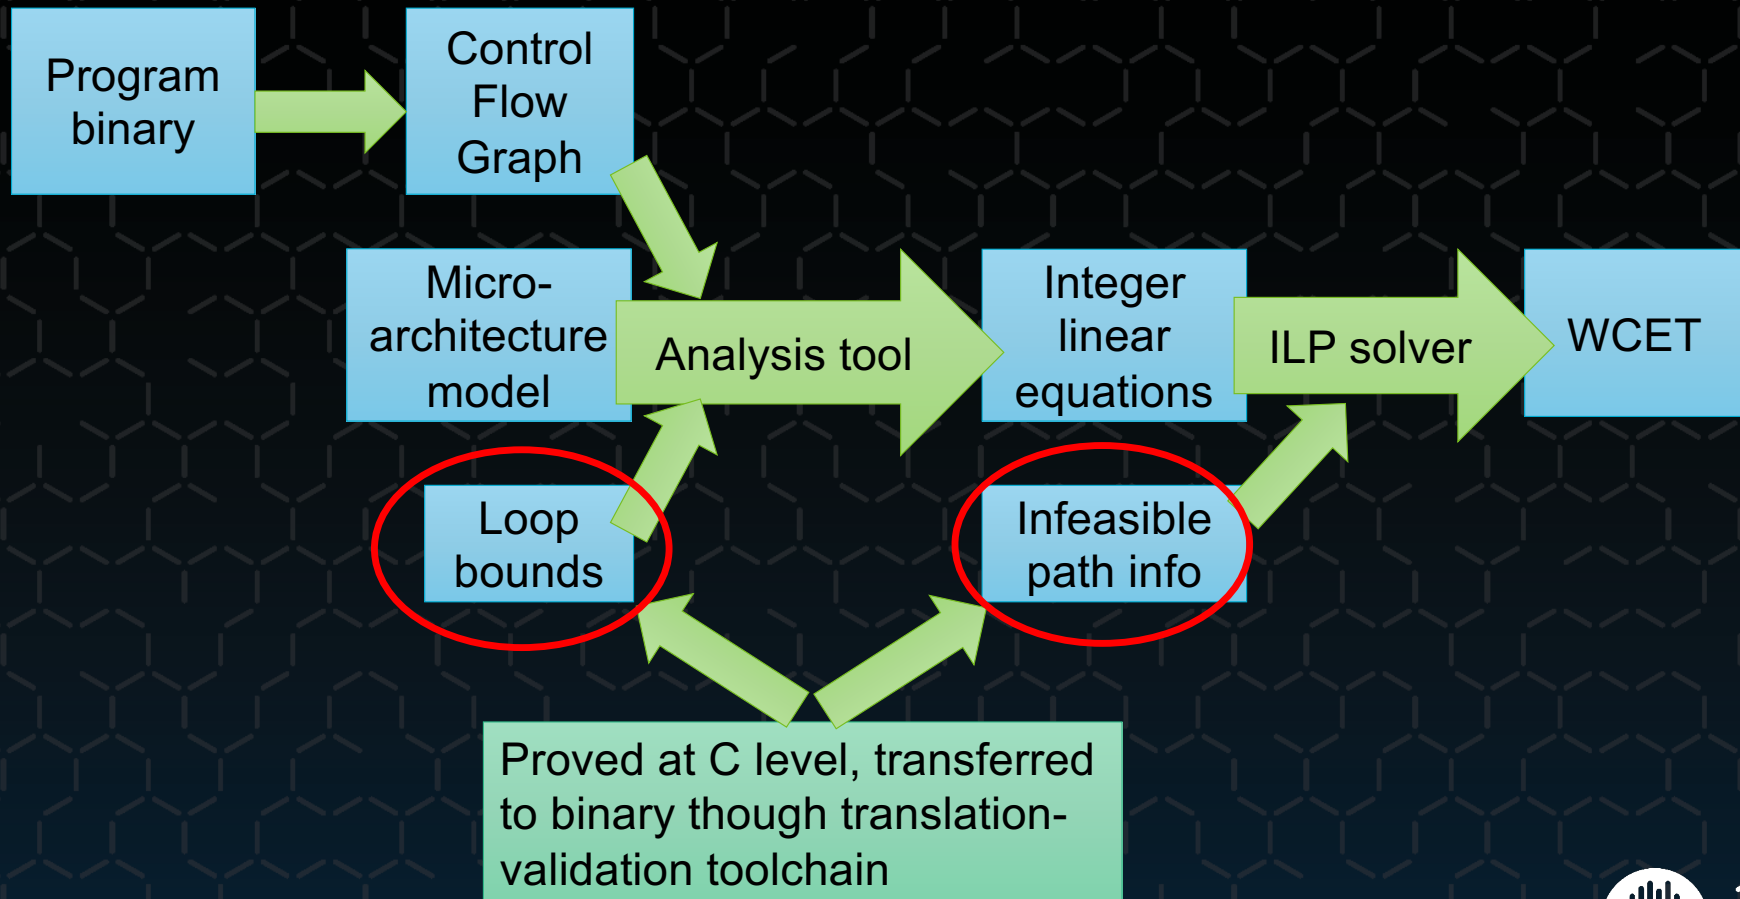
Temporal Isolation: WCET Analysis



se14



High-Assurance WCET Analysis



Temporal Isolation: Controlling Time



se14



Mixed Criticality: Critical + Untrusted



NW driver must preempt control loop

- ... to avoid packet loss
- Driver must run at high prio
- Driver must be trusted not to monopolise CPU

Runs every 100 ms
for few milliseconds

Sensor
readings

Critical:
Control
loop

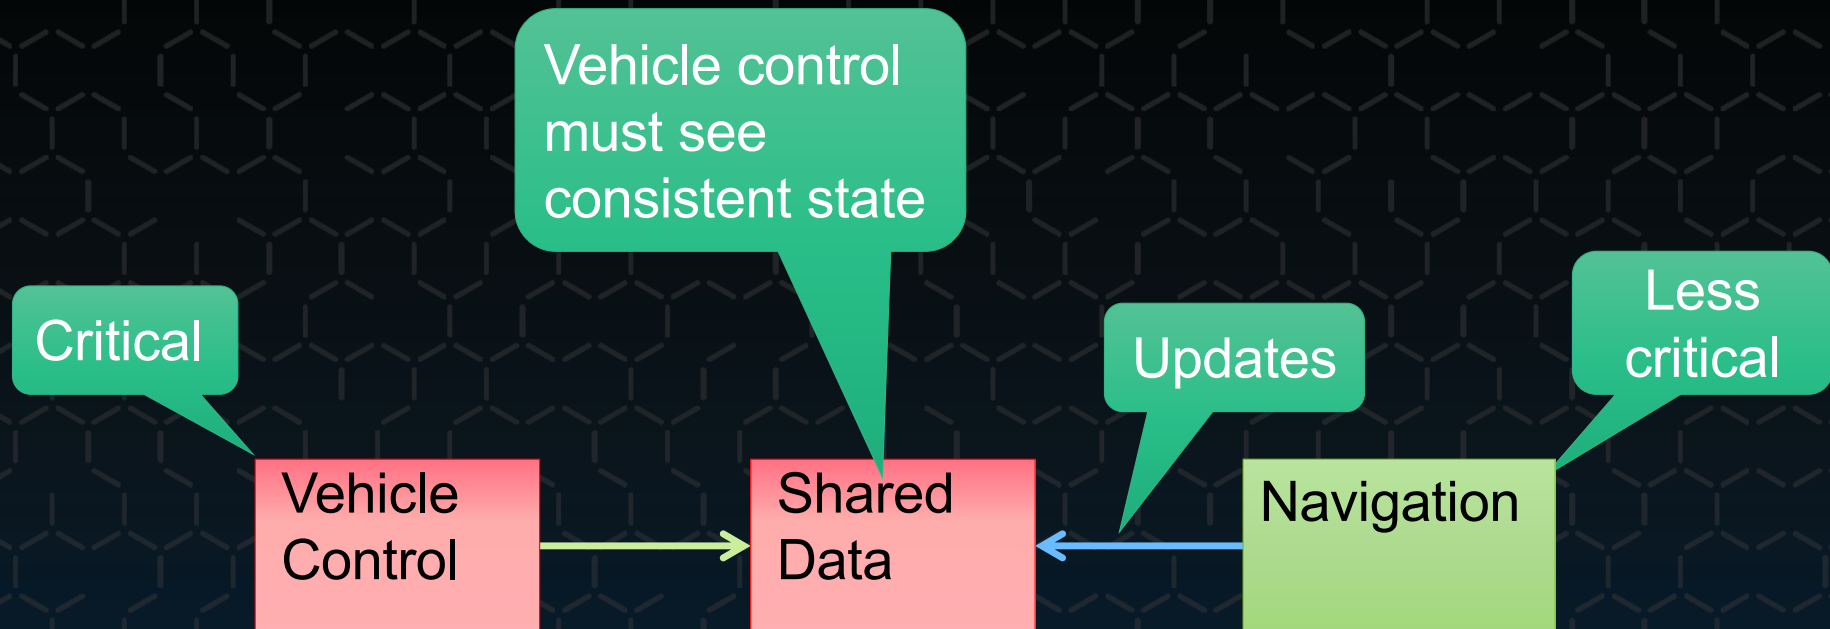


Untrusted:
NW
driver

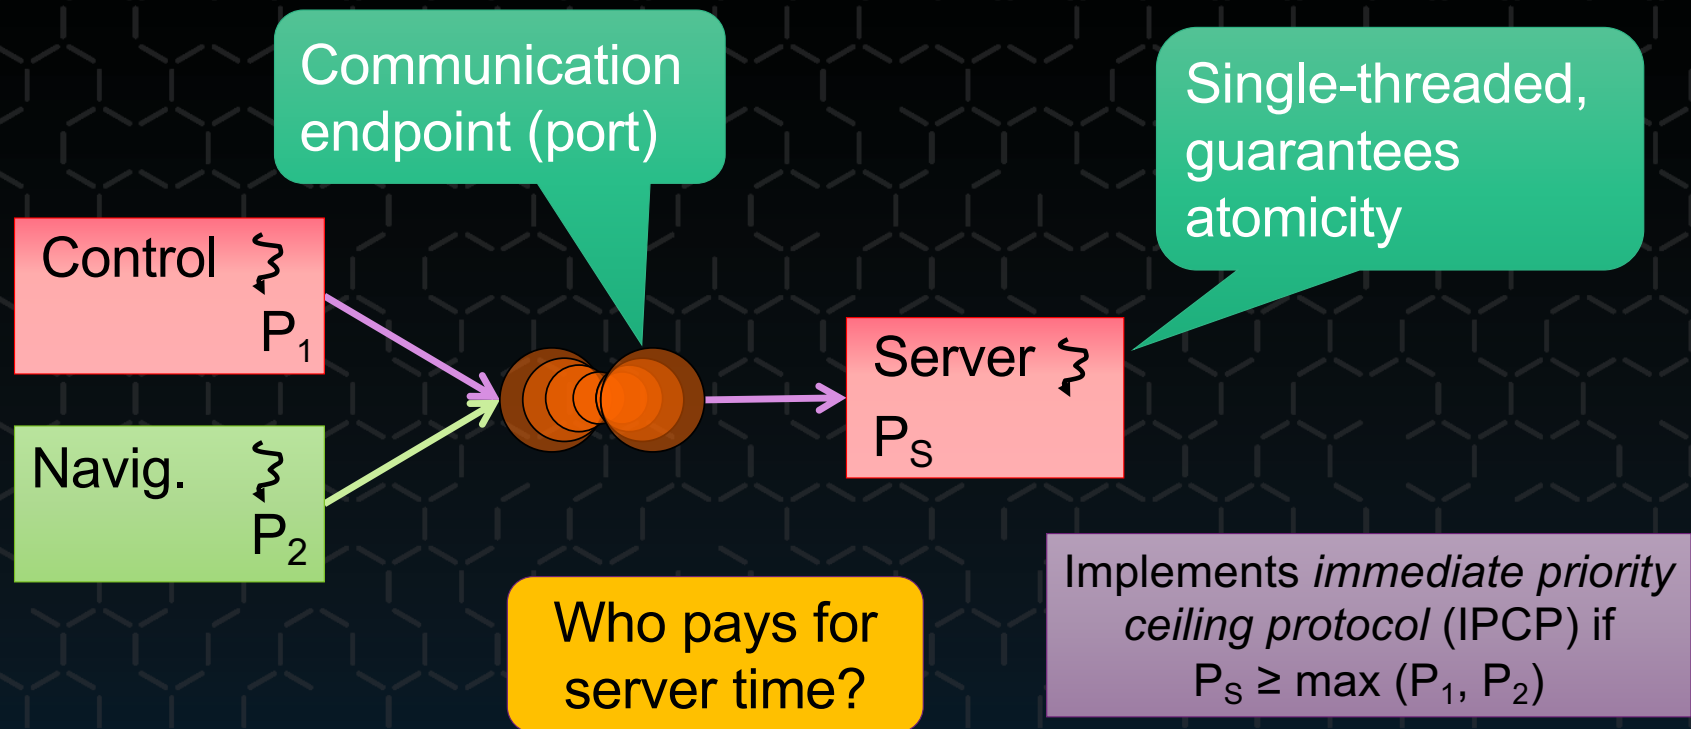
Runs frequently but for
short time (order of μ s)

NW
interrupts

MCS Challenge: Sharing



Sharing: Delegation to *Resource Server*



Solution: Time Capabilities



Classical thread attributes

- Priority
- Time slice

Not runnable
if null

Limits CPU
access –
sporadic server

Scheduling context object

- T: period
- C: budget ($\leq T$)

C = 2

T = 3



New thread attributes

- Priority
- Scheduling context capability

Capability
for time


Enables reasoning about
time and temporal isolation
for mixed-criticality systems


MCS with Scheduling Contexts



Runs every 100 ms
for few milliseconds

Sensor
readings

Control
loop 
P = low

NW
driver 
P = high


Runs frequently but for
short time (order of μ s)

NW
interrupts

C = 25,000

T = 100,000


Utilisation = 25%



C = 2

T = 3

Utilisation = 67%



Shared Server Time Charged to Client




Client is charged for server's time

Running



Running



Server runs on client's scheduling context

Timeout exception to deal with budget exhaustion

seL4 MCS Support



- **Time as a first-class resource:**

- Enforcement of delegatable time budgets
- Suitable for formal reasoning
- Verification to be completed this year

- **Status:**

- Functional correctness of MCS extensions presently being verified for Arm and RISC-V

- **To Do:**

- Proving scheduler properties
- Formal framework for reasoning about timeliness of applications



Thank You!

