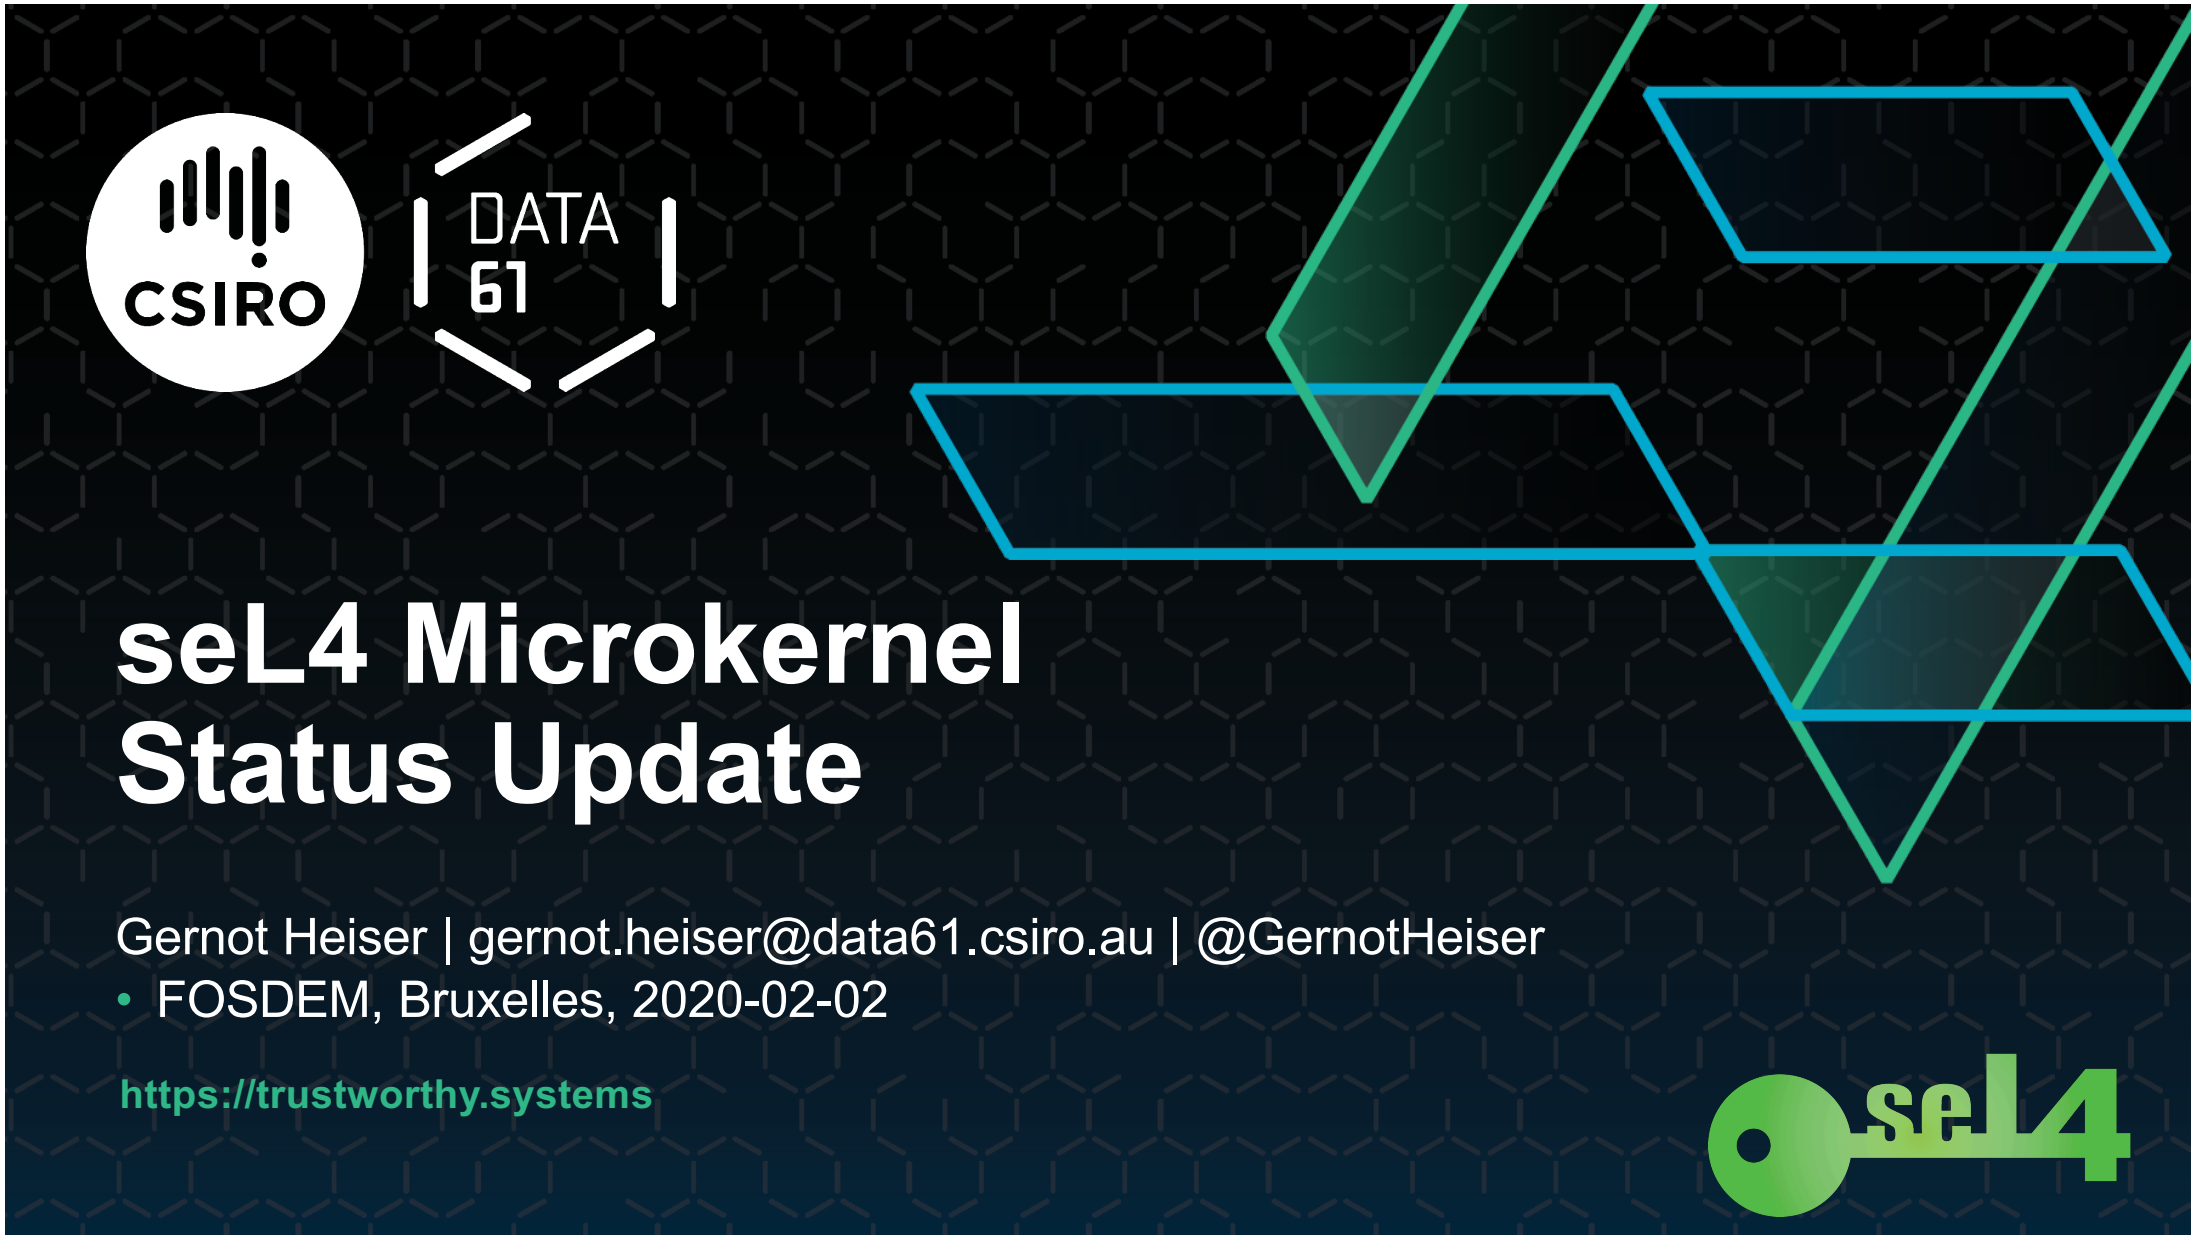

seL4 Microkernel Status Update

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

- FOSDEM, Bruxelles, 2020-02-02

https://trustworthy.systems

What is seL4?

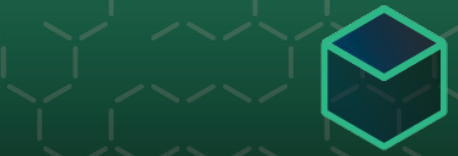

CSIRO DATA 61

seL4: Assurance and Performance

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

CSIRO DATA 61

World's Most Secure OS: Arm v7

Confidentiality

Integrity

Availability

Proof

Proof

Proof

Model enforces security

Abstract Model

Functional correctness: C code only behaves as specified

Proof

Translation validation: Binary retains C-code semantics

C Imple-mentation

Proof

Sound worst-case execution time bound

Binary code

Limitations (work in progress):

- Kernel initialisation not yet verified
- MMU & caches modelled abstractly
- Timing channels not ruled out

CSIRO DATA 61

Military-Strength Security

Unmanned Little Bird (ULB)

DARPA HACMS: Retrofit existing system!

Autonomous trucks

Secure Comms Dongle

Cross-Domain Desktop Compositor

seL4 on RISC-V

CSIRO DATA 61

Background: HENSOLD Cyber

Disclosure: I have an interest in HENSOLDT Cyber

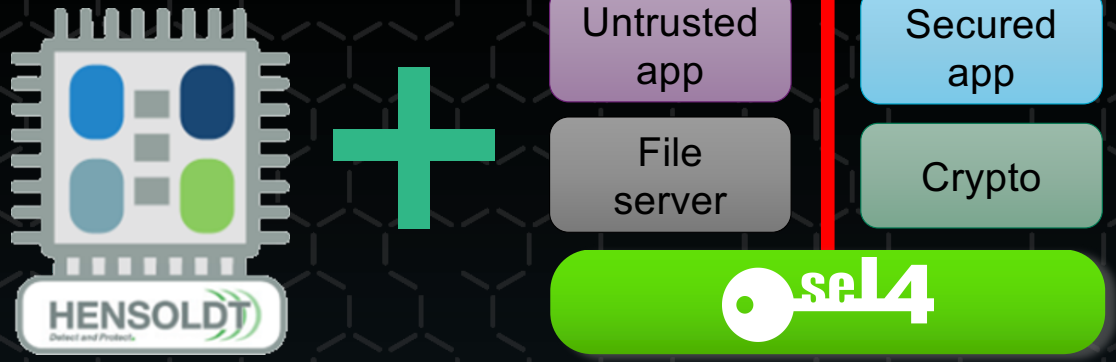

Munich-based startup

- Secure RISC-V processor
- Based on open-source Ariane core (ETH)
- Supply chain secured through logic encryption
- Secure OS based on seL4
- Targets defence, industrial control, critint, automotive

Performance on RV64

Message-passing round-trip latency in cycles

Not yet fully optimised!

Arch	x86 32b	x86 64b	Arm 32b	Arm 64b	RISC-V 64b
Intra address space	427	565	625	752	690
Inter address space	752	1041	625	752	1006

Meltdown-workaround disabled (else much slower!)

No ASIDS on **HiFive Unleashed**, else inter-AS would be same as intra-AS

Hypervisor extensions supported in branch, tracking draft spec

Verification: RISC-V Status

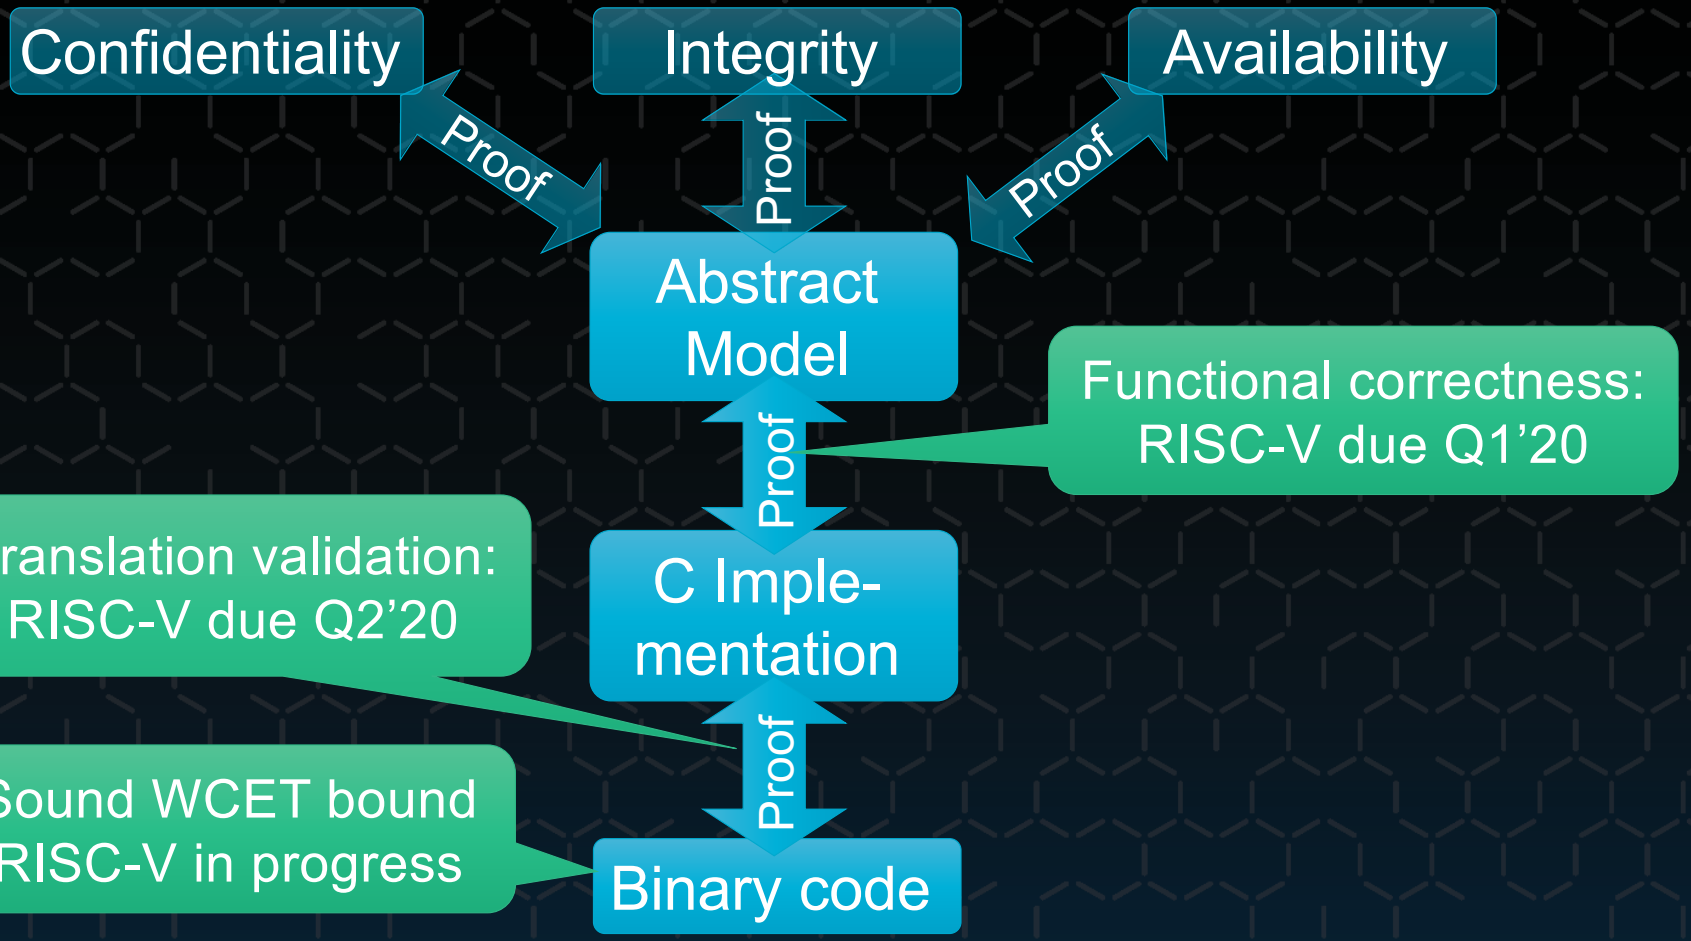

Experience with RISC-V Architecture

- Kernel port straightforward:
  - simple and clean RISC architecture
- Verification benefitted from cleanness
  - … but some challenges from less typing in page tables
- Hypervisor (draft) extensions even simpler
- M (machine) mode makes firmware explicit
  - configures HW, delegates to S (supervisor) mode
  - emulates features not implemented in HW
  - should be verified
- Extensibility of ISA could be a concern
  - could undermine portability
- Formal ISA spec is great!

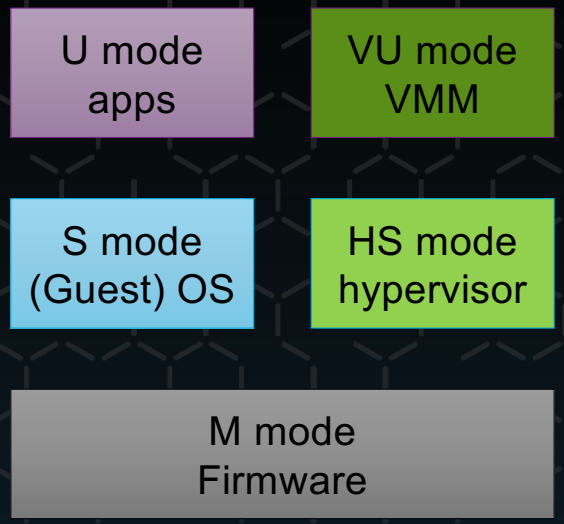

Mixed-Criticality Scheduling

(FOSDEM'19 Refresher)

CSIRO DATA 61

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

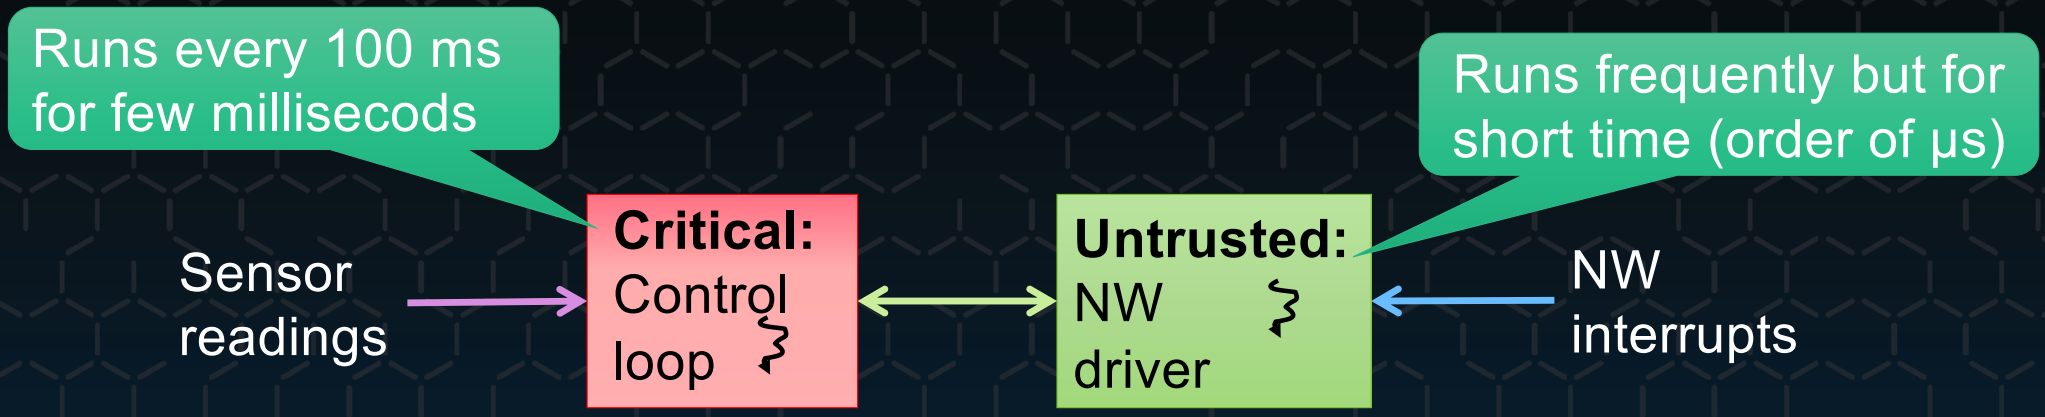

CSIRO DATA 61

MCS Challenge: Sharing

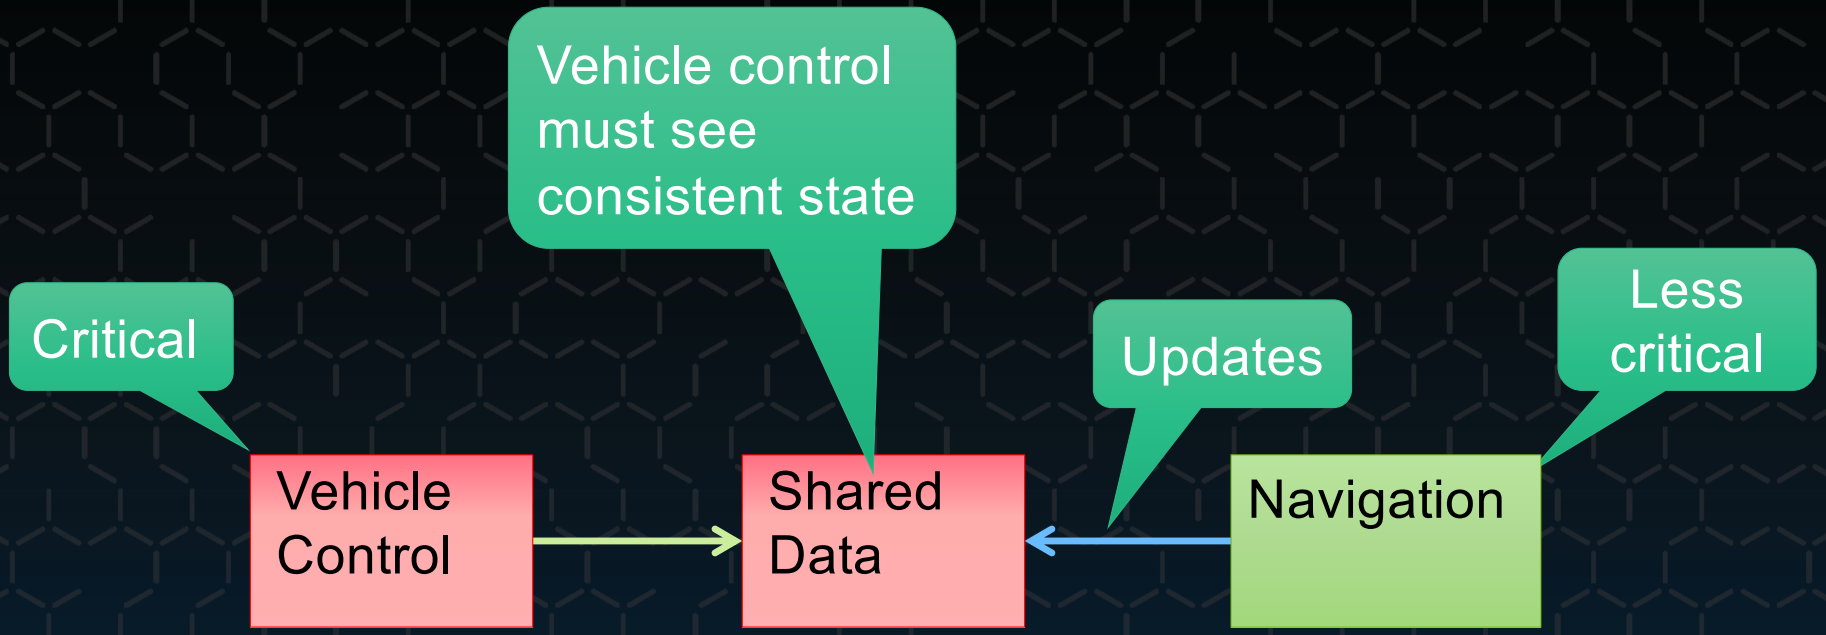

Sharing Through *Resource Server*

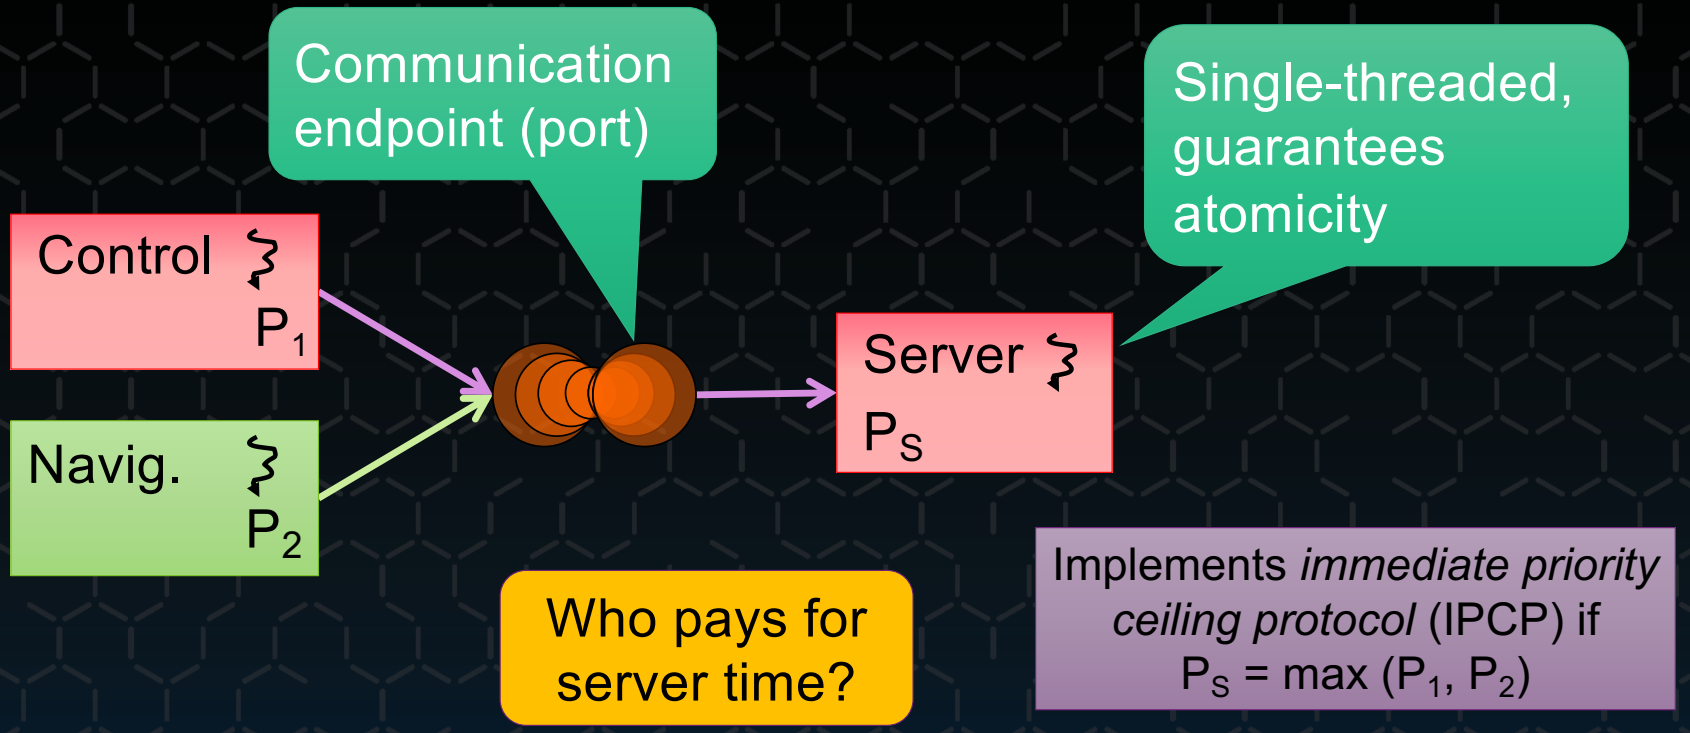

Solution: Time Capabilities

Classical thread attributes

- Priority
- Time slice

Not runnable if null

New thread attributes

- Priority
- Scheduling context capability

Capability for time

Limits CPU access!

Scheduling context object
- T: period
- C: budget (≤ T)

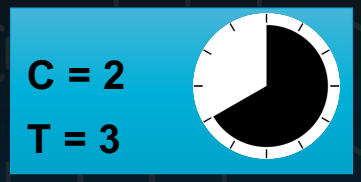

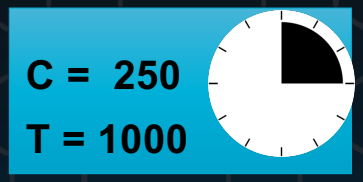

Enables reasoning about time and temporal isolation for mixed-criticality systems

CSIRO DATA 61

Time Caps (MCS) Kernel Verification

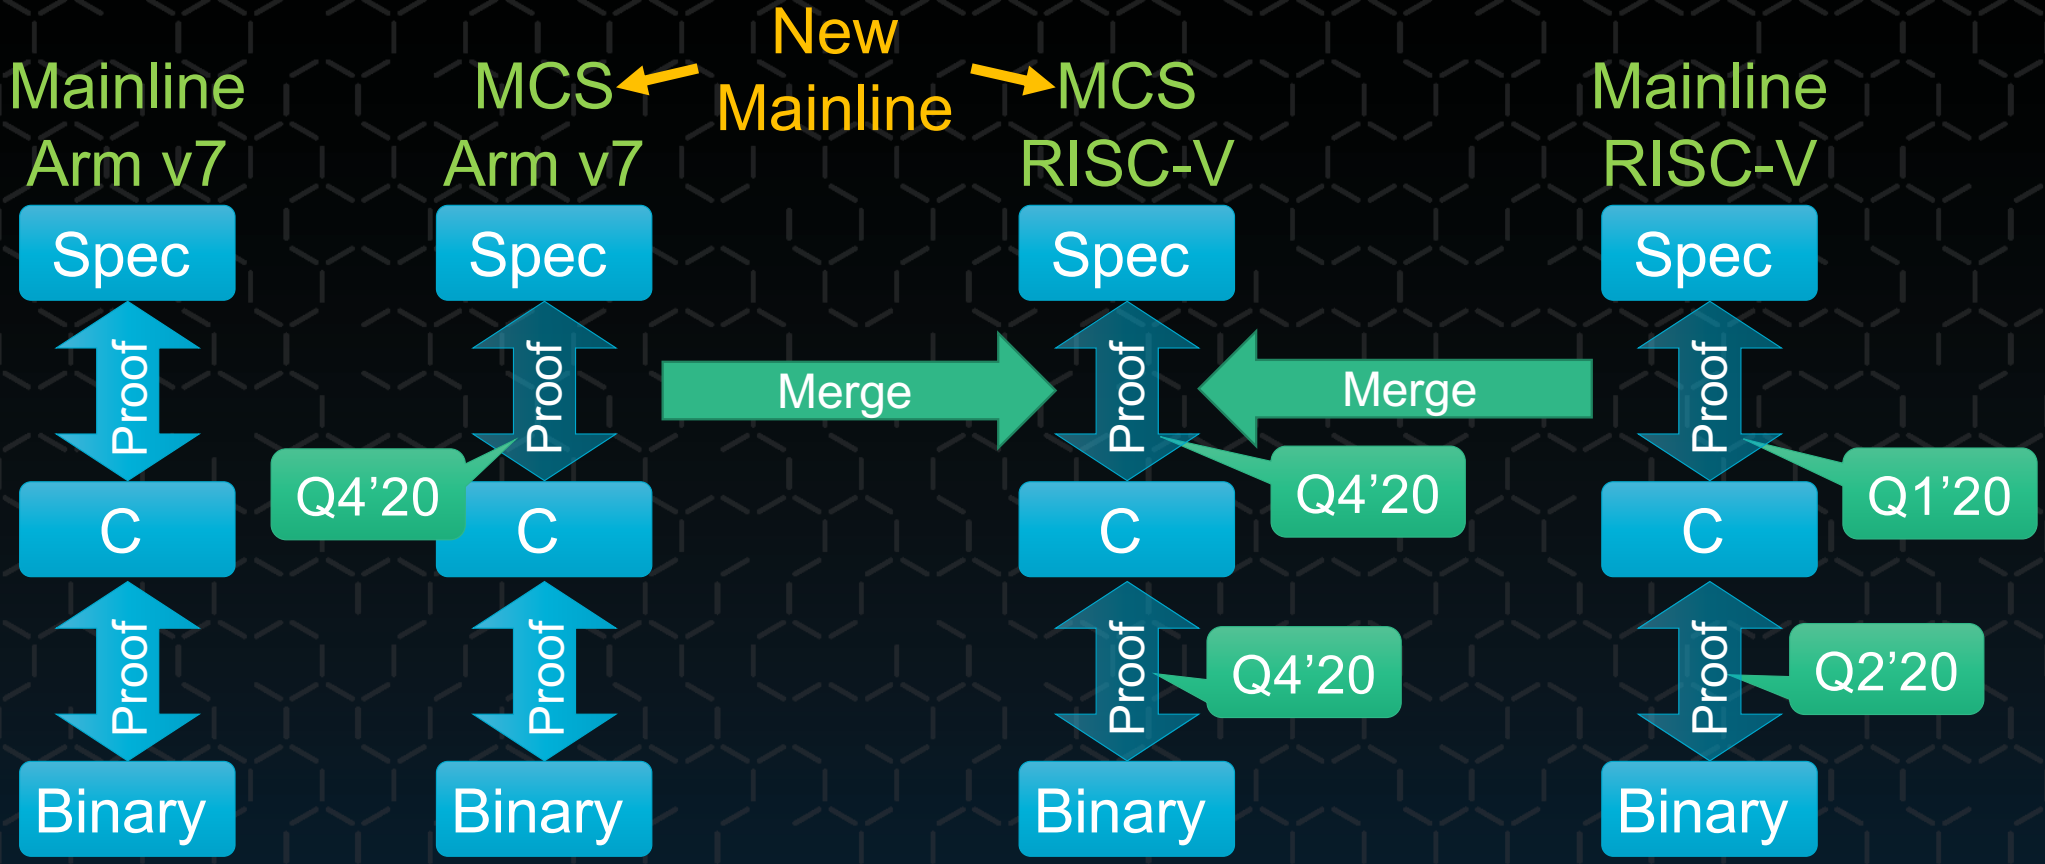

Community/ Ecosystem

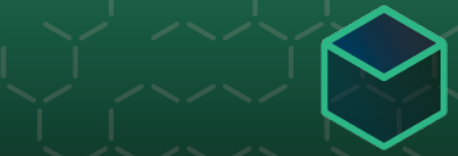

Experience with RISC-V Foundation

Security Standing Committee

- Invited me on
- Very receptive and supportive
- Committed to making RISC-V "most secure architcture"
- Facilitated engagement with Privspec TC (now Standing Committee)

Privileged Spec Tech Committee

- Hypervisor-extension feedback well received
  - Easy engagement
  - Constructive proposal from TC chair addressing our issues
- Time-protection slow to get traction
  - Now good engagement, hopefully progress soon

- Open but skeptical
- They need to manage conflicting ideas
- Keen to get "most secure arch" recognition

We Are Creating the seL4 Foundation!

Aims:

- Provide a neutral entity for coordinating & enhancing seL4 ecosystem
- Grow adoption of seL4
- Improve (organisational and individual) community participation & cooperation
  - Developers
  - Adopters
- Develop / standardise seL4 system
  - kernel & proofs
  - libraries, services, tools
- Protect and promote the seL4 brand
  - prevent reputational damage from using modified seL4 (verification invalidated)
- Provide platform for pooling funds for critical “big-ticket” items (verification)

CSIRO DATA 61

Foundation Structure

seL4 Foundation

seL4 Board

seL4 Fund Charter

seL4 Directed Fund $$

LF Projects LLC

seL4 Series LLC

seL4 TM

https://sel4.systems

seL4 Technical Charter

Technical Project

Contributor

Membership and Governance

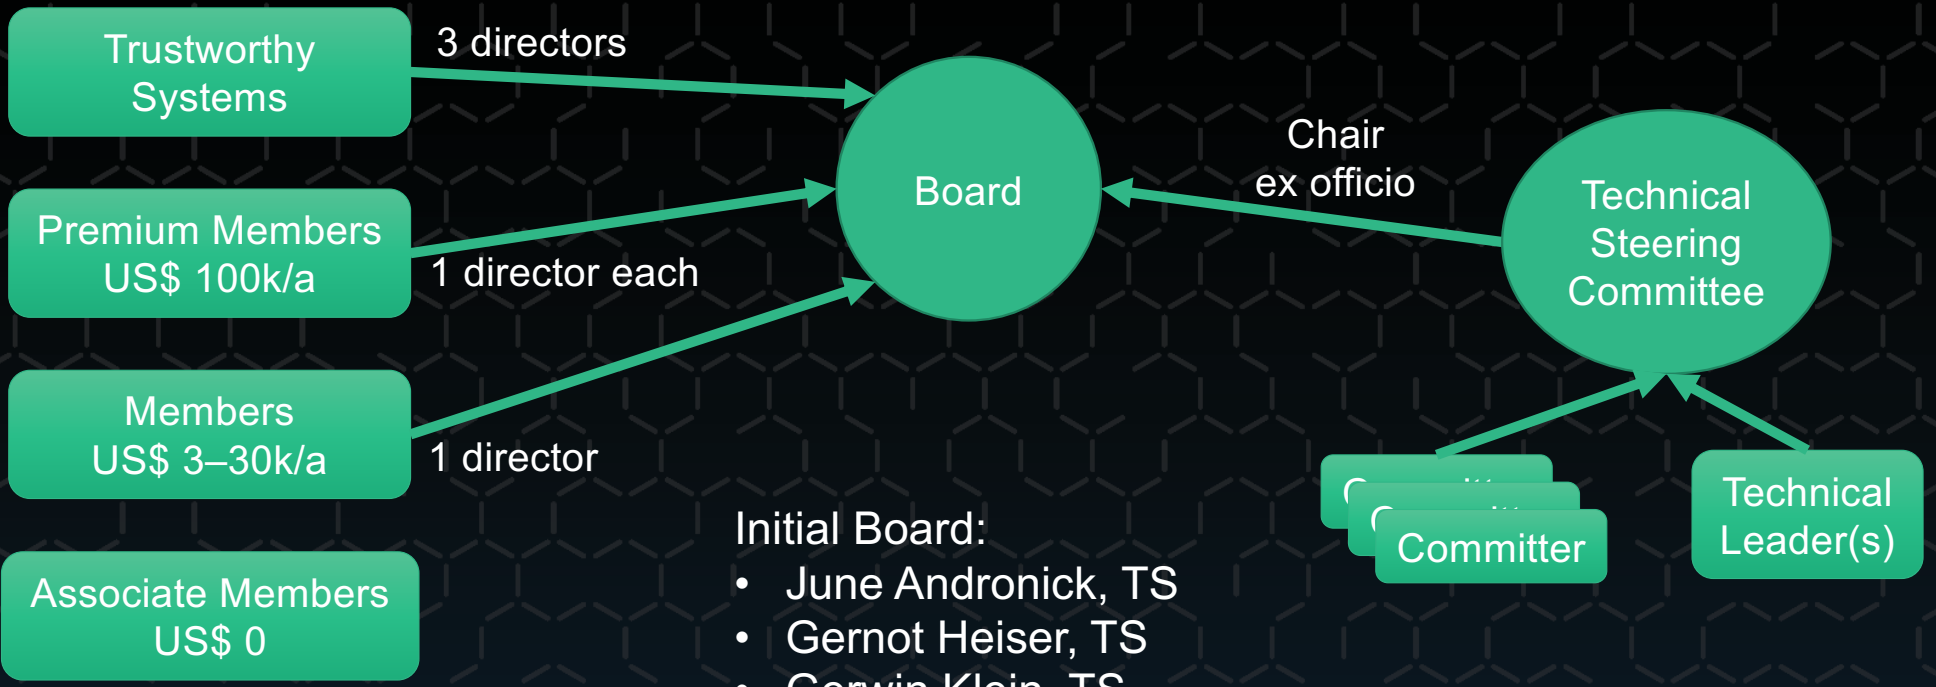

Associate Members
US$ 0

Note: members must be financial members of Linux Foundation!

Initial Board:

- June Andronick, TS
- Gernot Heiser, TS
- Gerwin Klein, TS
- John Launchbury, Galois (ex DARPA)
- Sascha Kegreiß, HENSOLDT Cyber
- Daniel Potts, Ghost Locomotion

Community Engagement

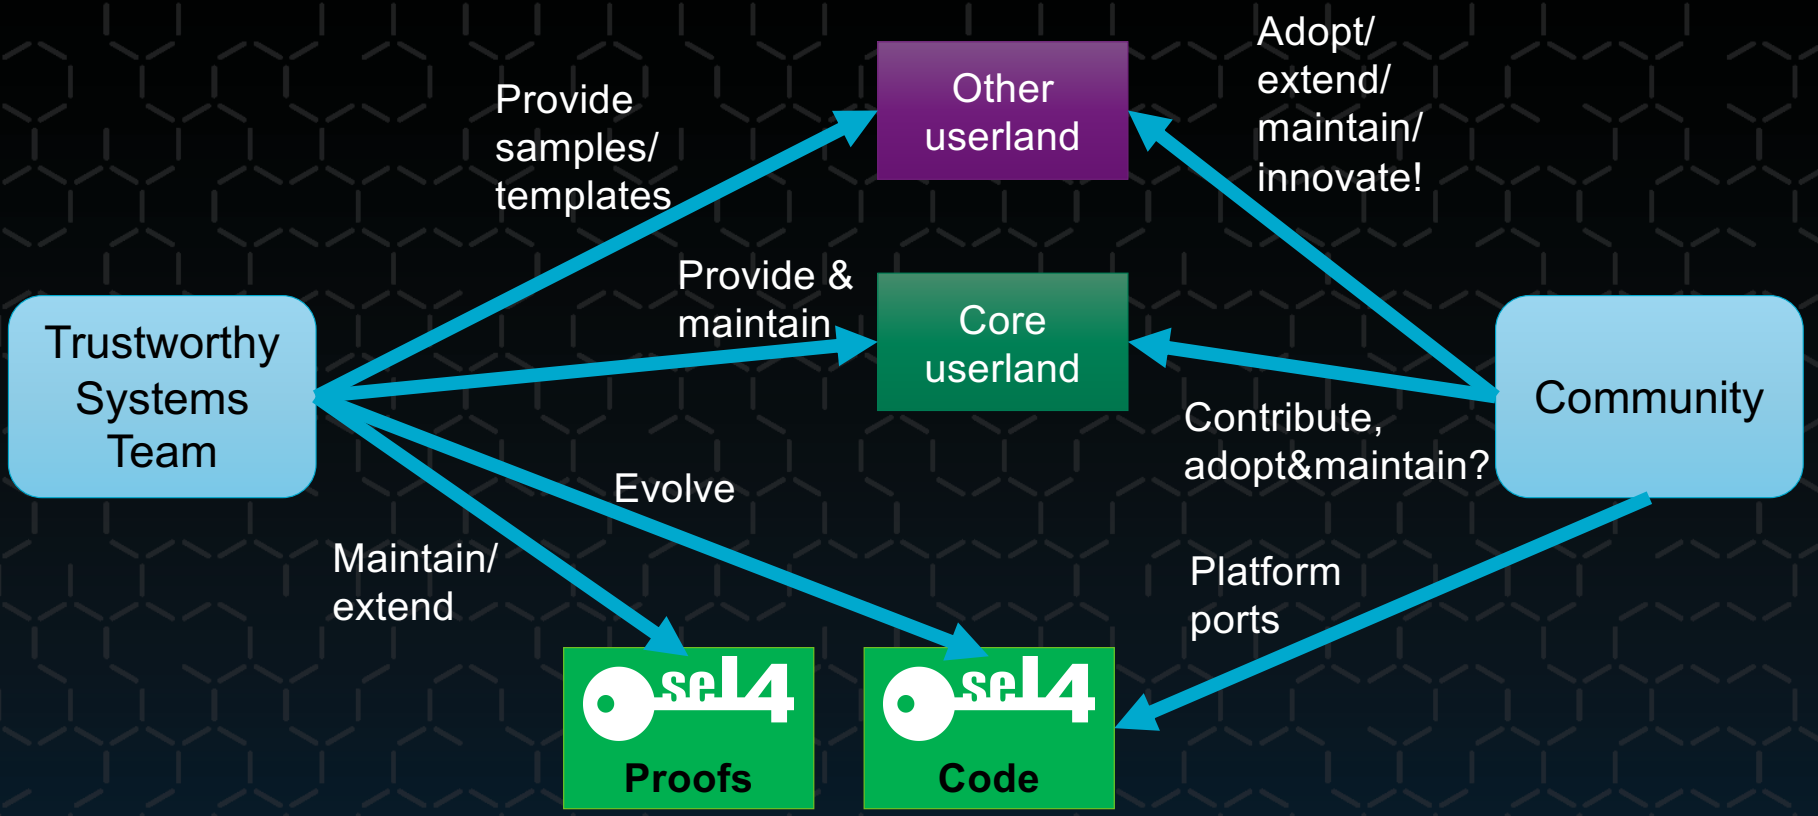

Foundation Status

- Legal docs (fund charter & technical charter) approved by Linux Foundation
- Trademark ready for transfer to Foundation
- Initial board appointed
- Interim web site shows structure, “Principles” and legal docs
- Hopefully days away from being able to set up members
  - Mail foundation@sel4.systems if you’re interested in joining!
  - Will make announcement on seL4.systems mailing lists

https://sel4.systems/Foundation