



Formally Verified System Initialisation

Andrew Boyton

June Andronick, Callum Bannister,
Matthew Fernandez, Xin Gao,
David Greenaway, Gerwin Klein,
Corey Lewis and Thomas Sewell



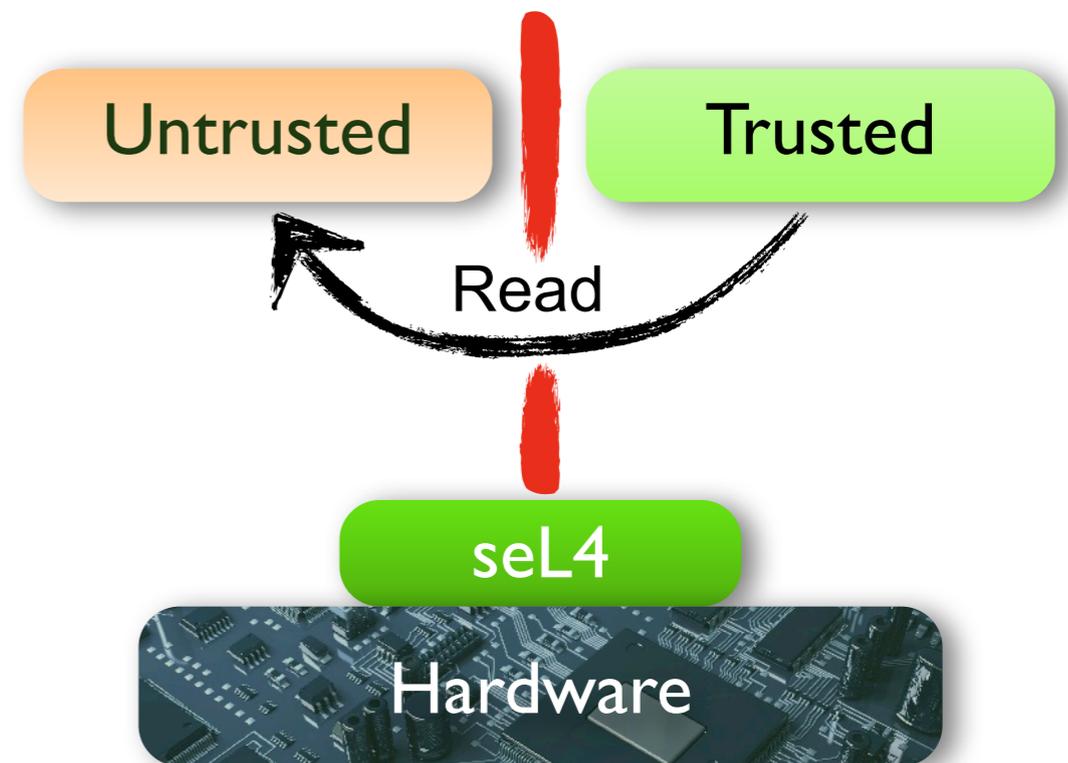
Australian Government
**Department of Broadband, Communications
and the Digital Economy**
Australian Research Council

NICTA Funding and Supporting Members and Partners



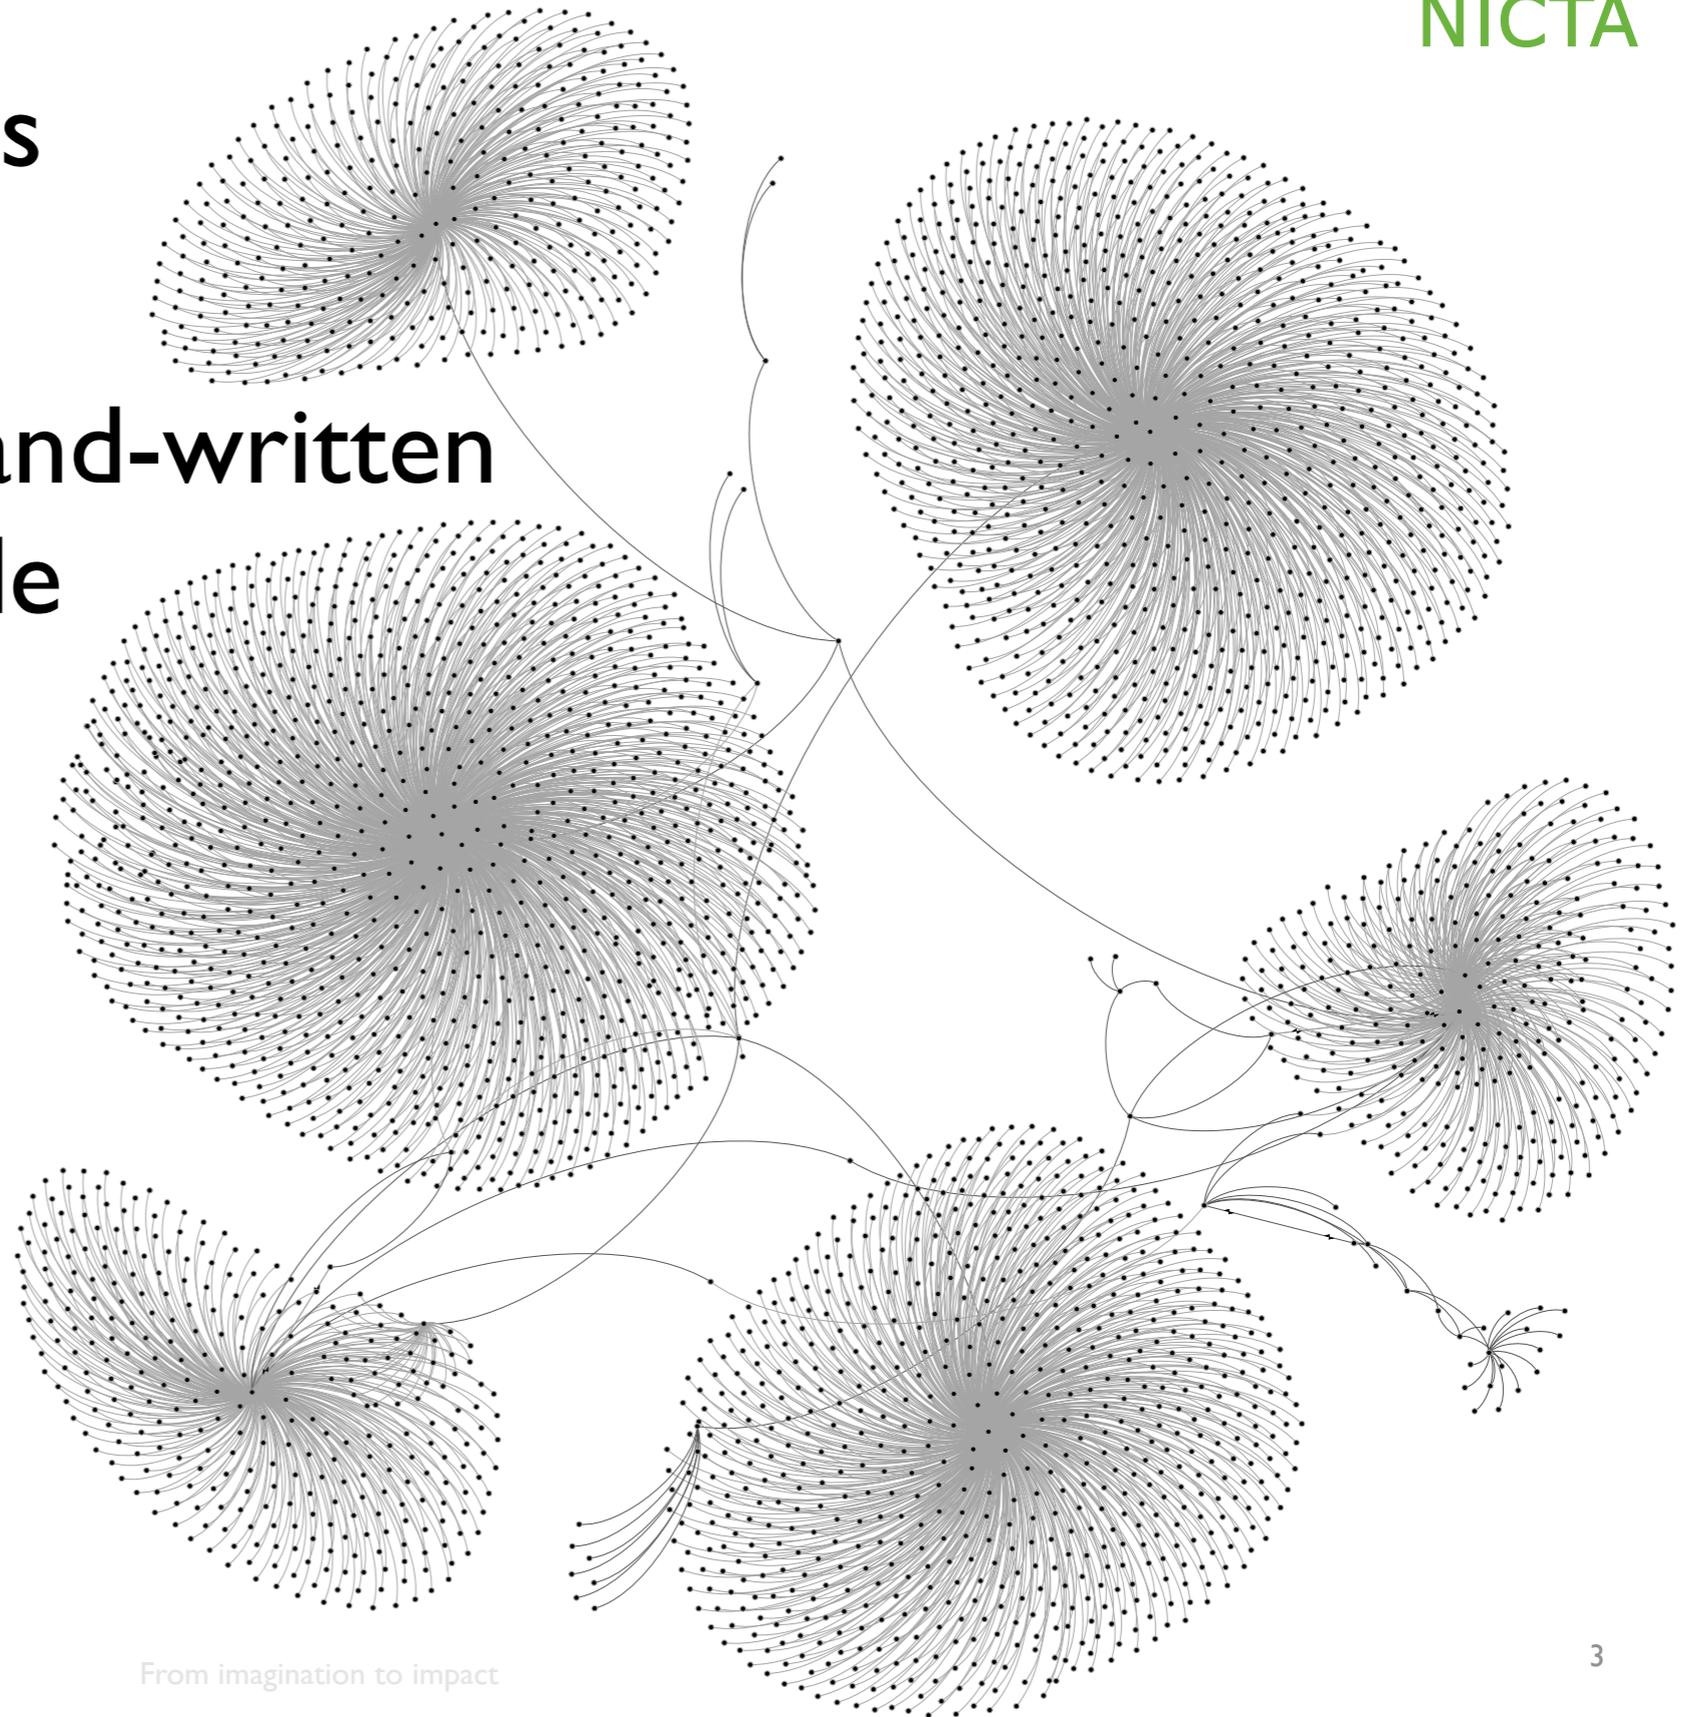
Motivation

- Create high-assurance systems
 - Pacemaker
 - Unmanned aerial vehicle (UAV)
 - Secure access terminal
- Capability-based system
 - seL4 microkernel
- Two problems:
 - Does our capability system correctly enforce the capabilities?
 - ***How do we set up a system with the desired capabilities?***



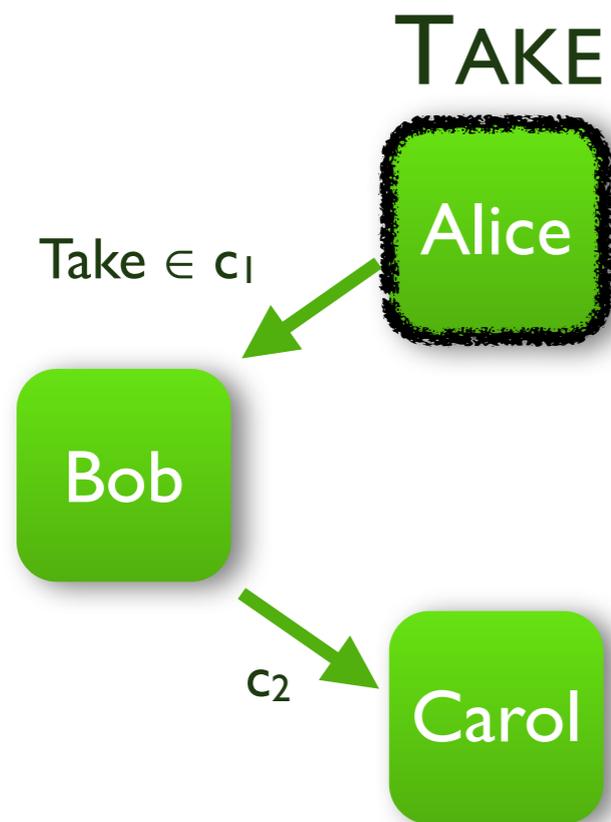
Initialisation code can get big

- 3,600 capabilities
- 3,500 objects
- 2,500 lines of hand-written initialisation code



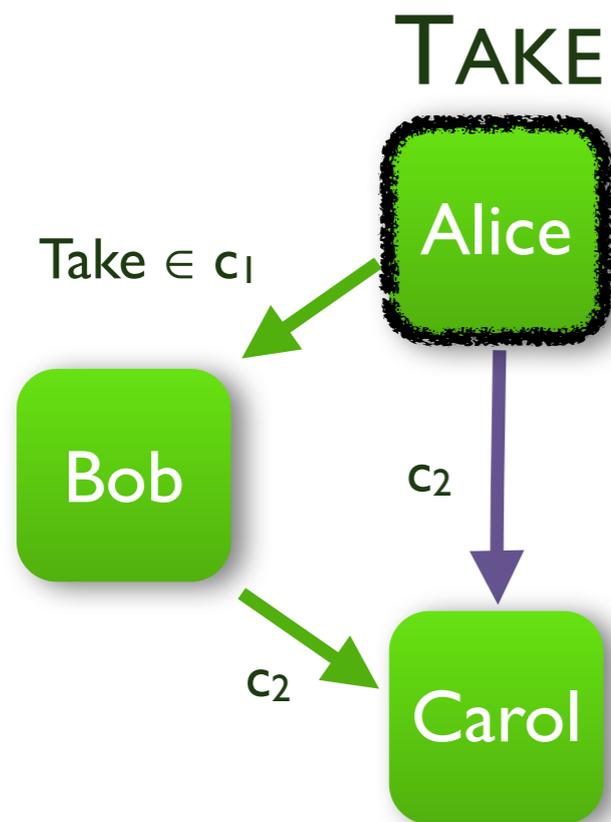
Capability systems

- Take-grant model



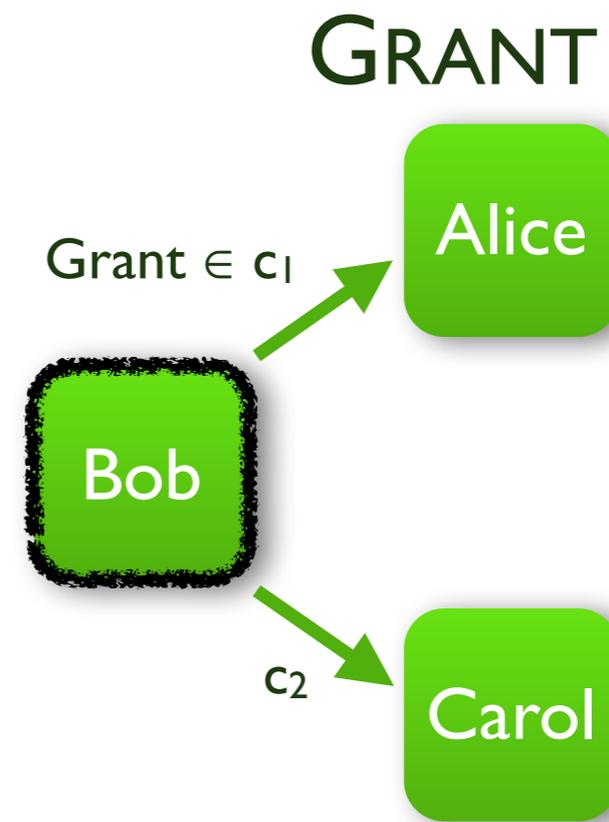
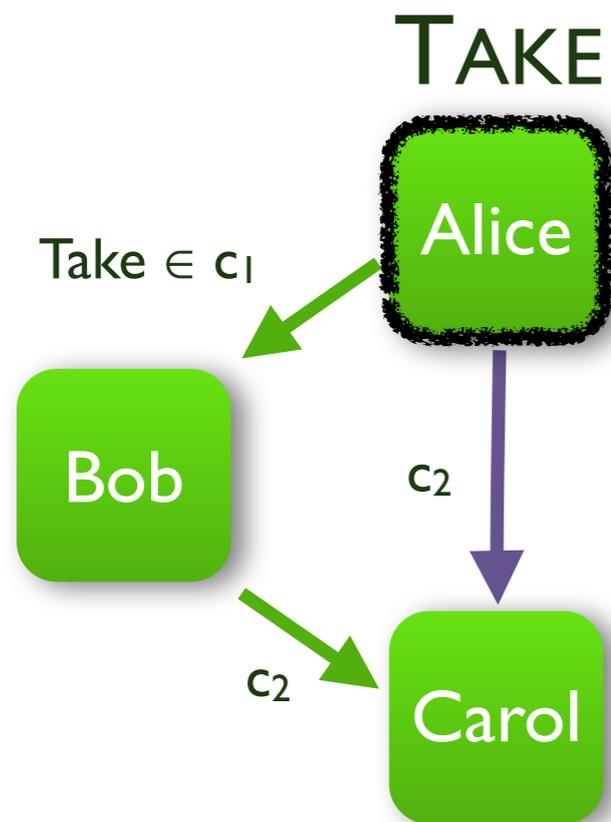
Capability systems

- Take-grant model



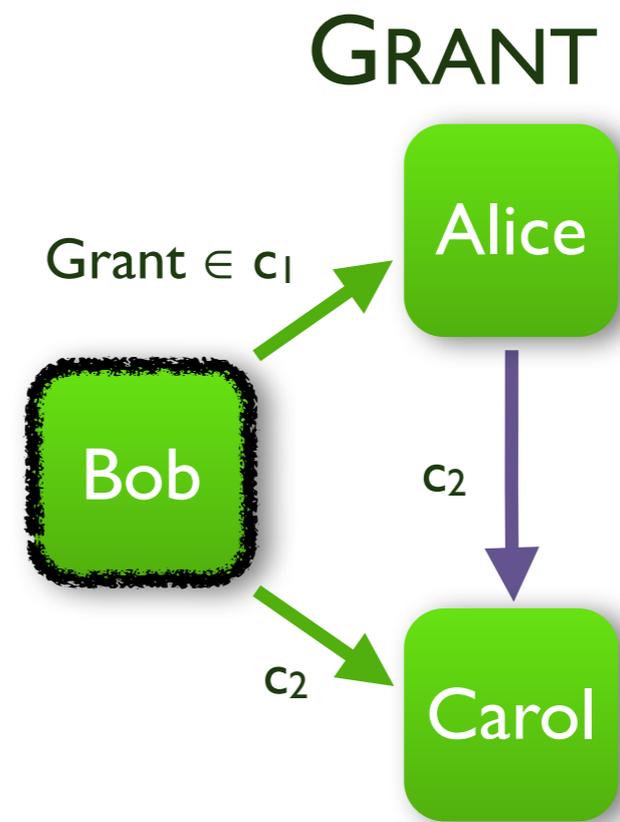
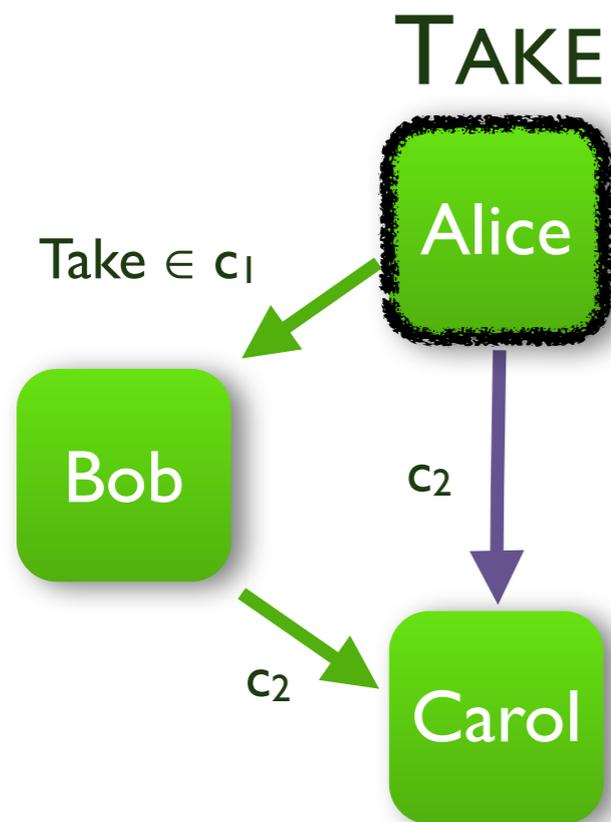
Capability systems

- Take-grant model

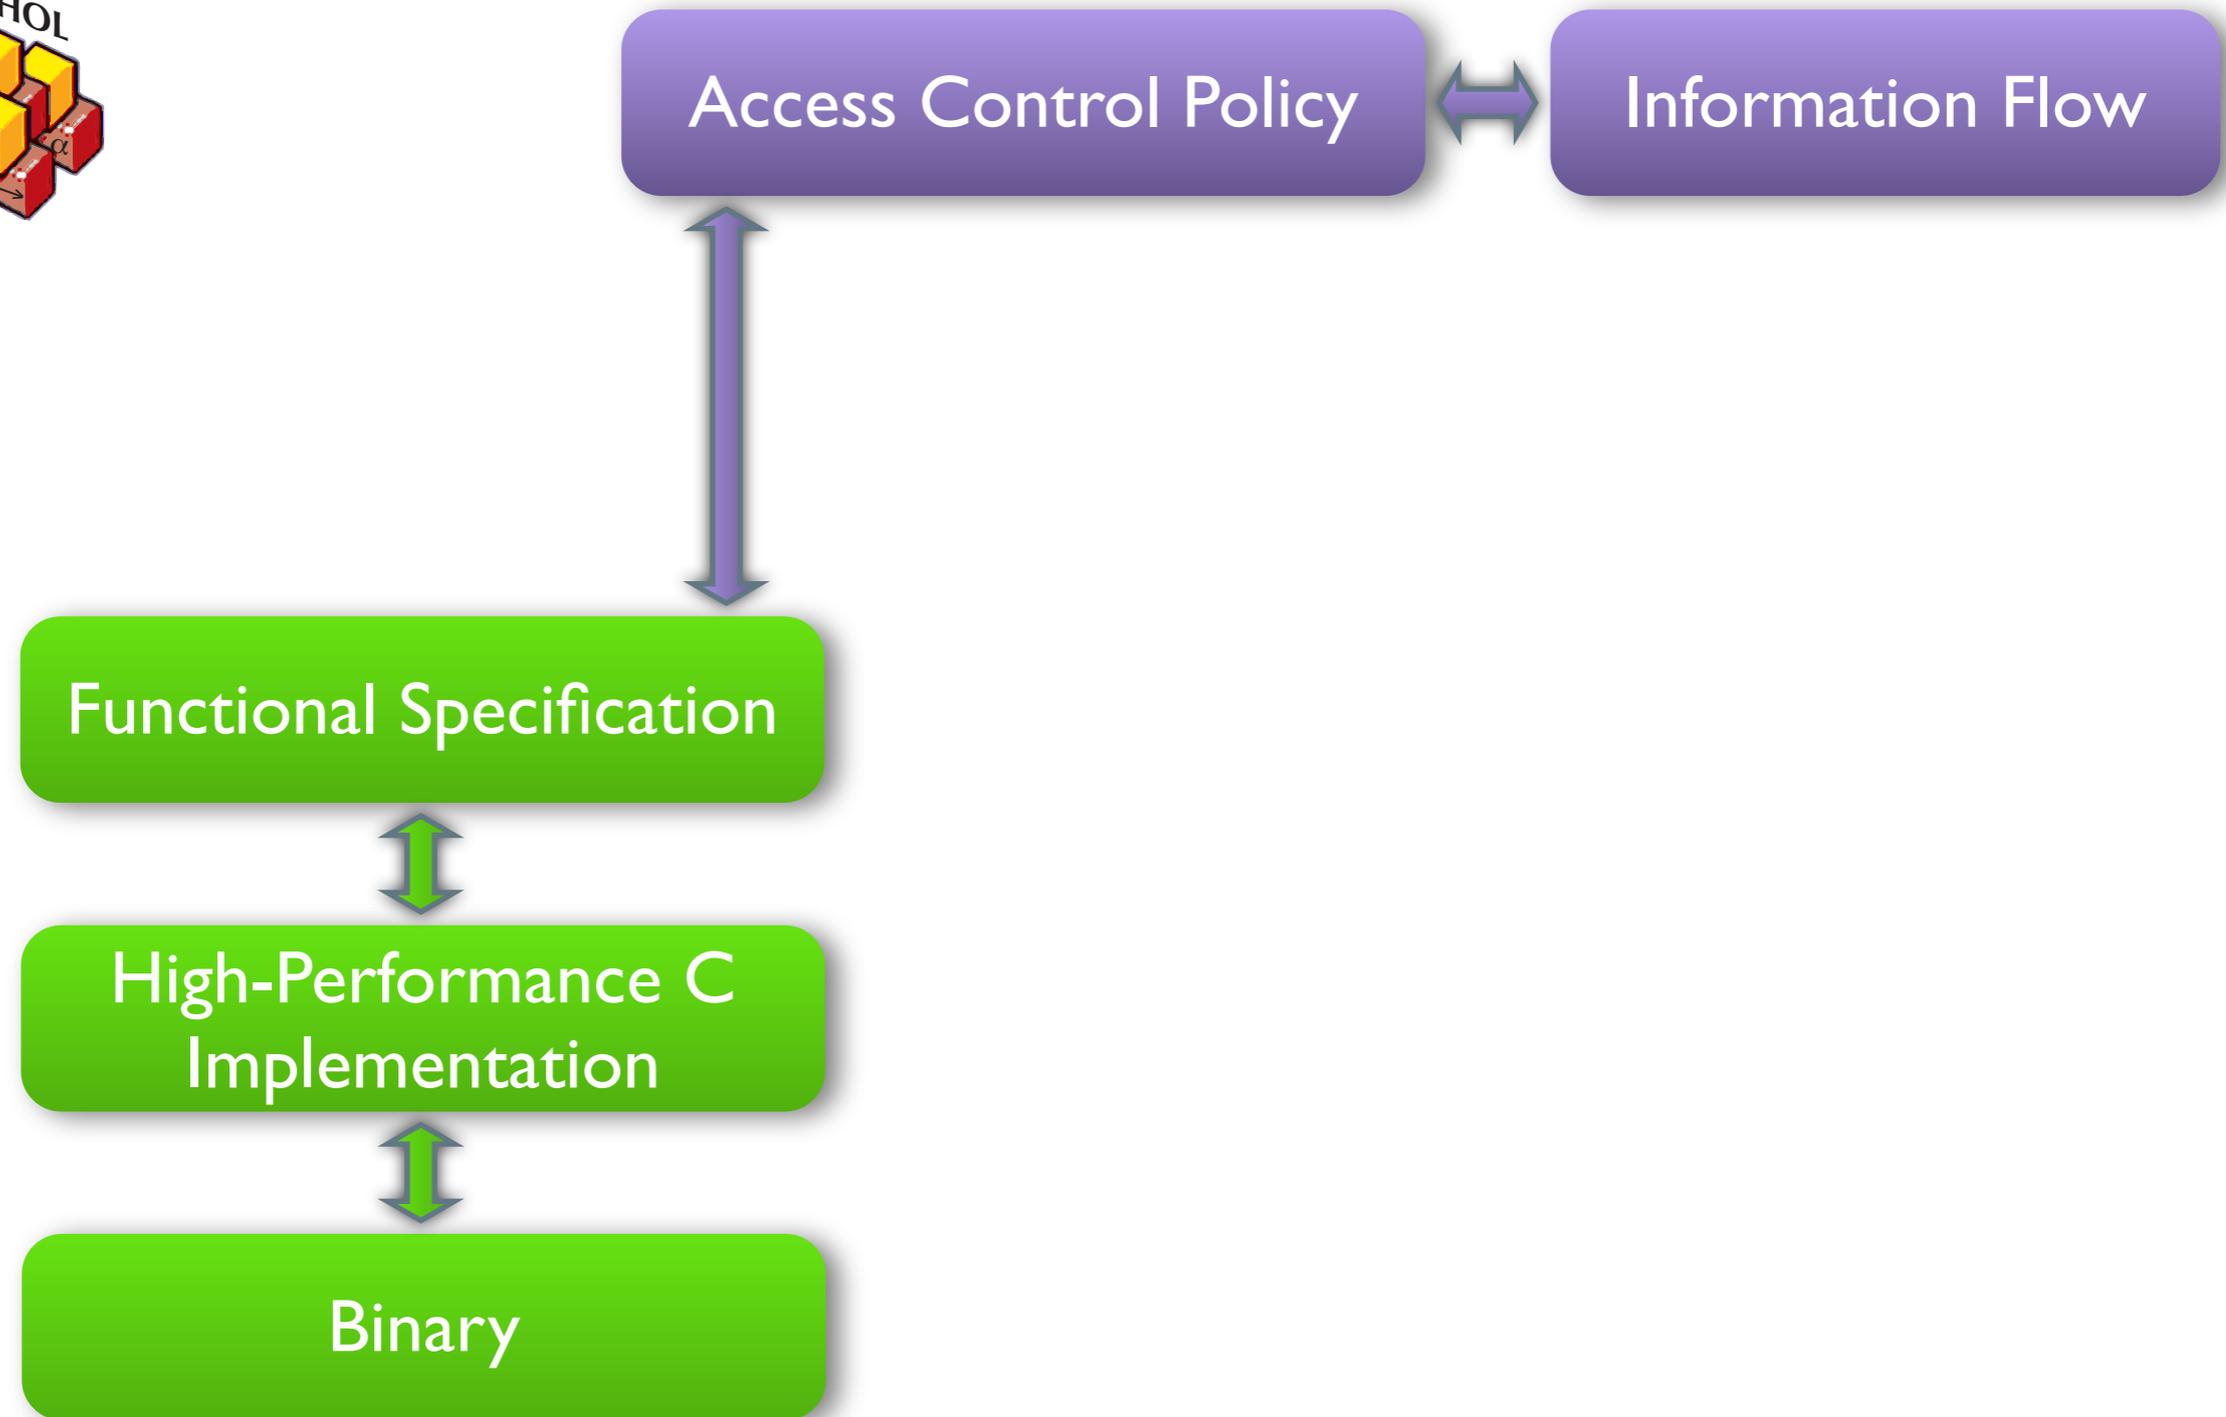


Capability systems

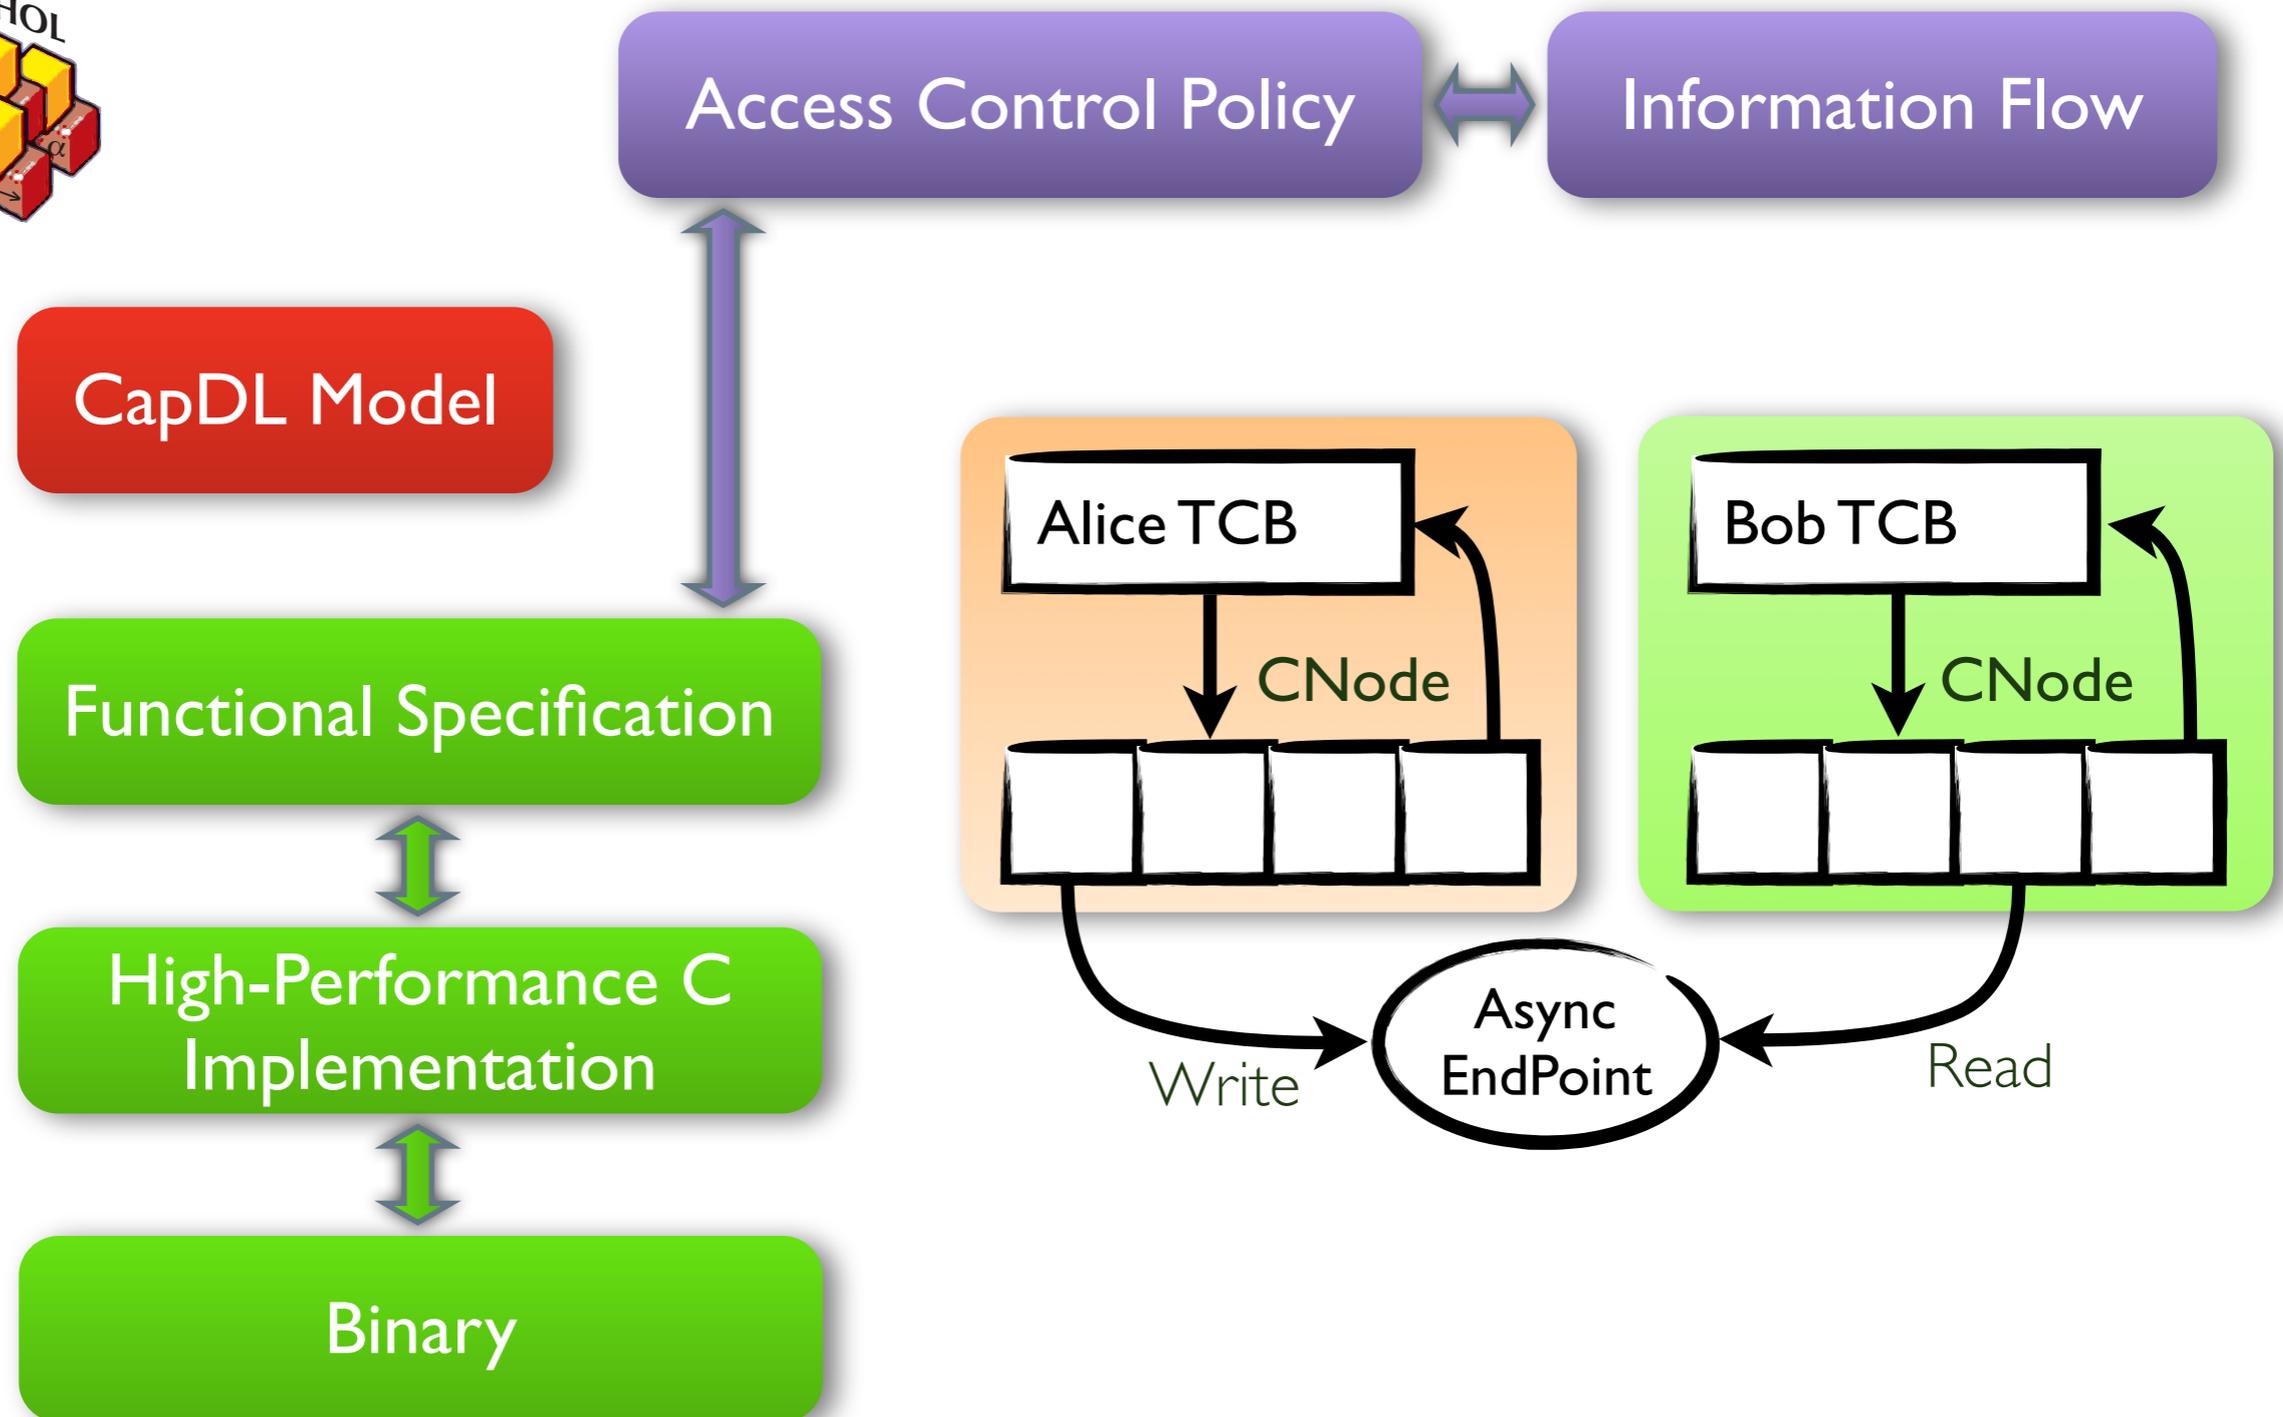
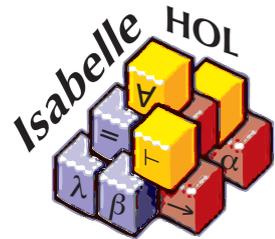
- Take-grant model



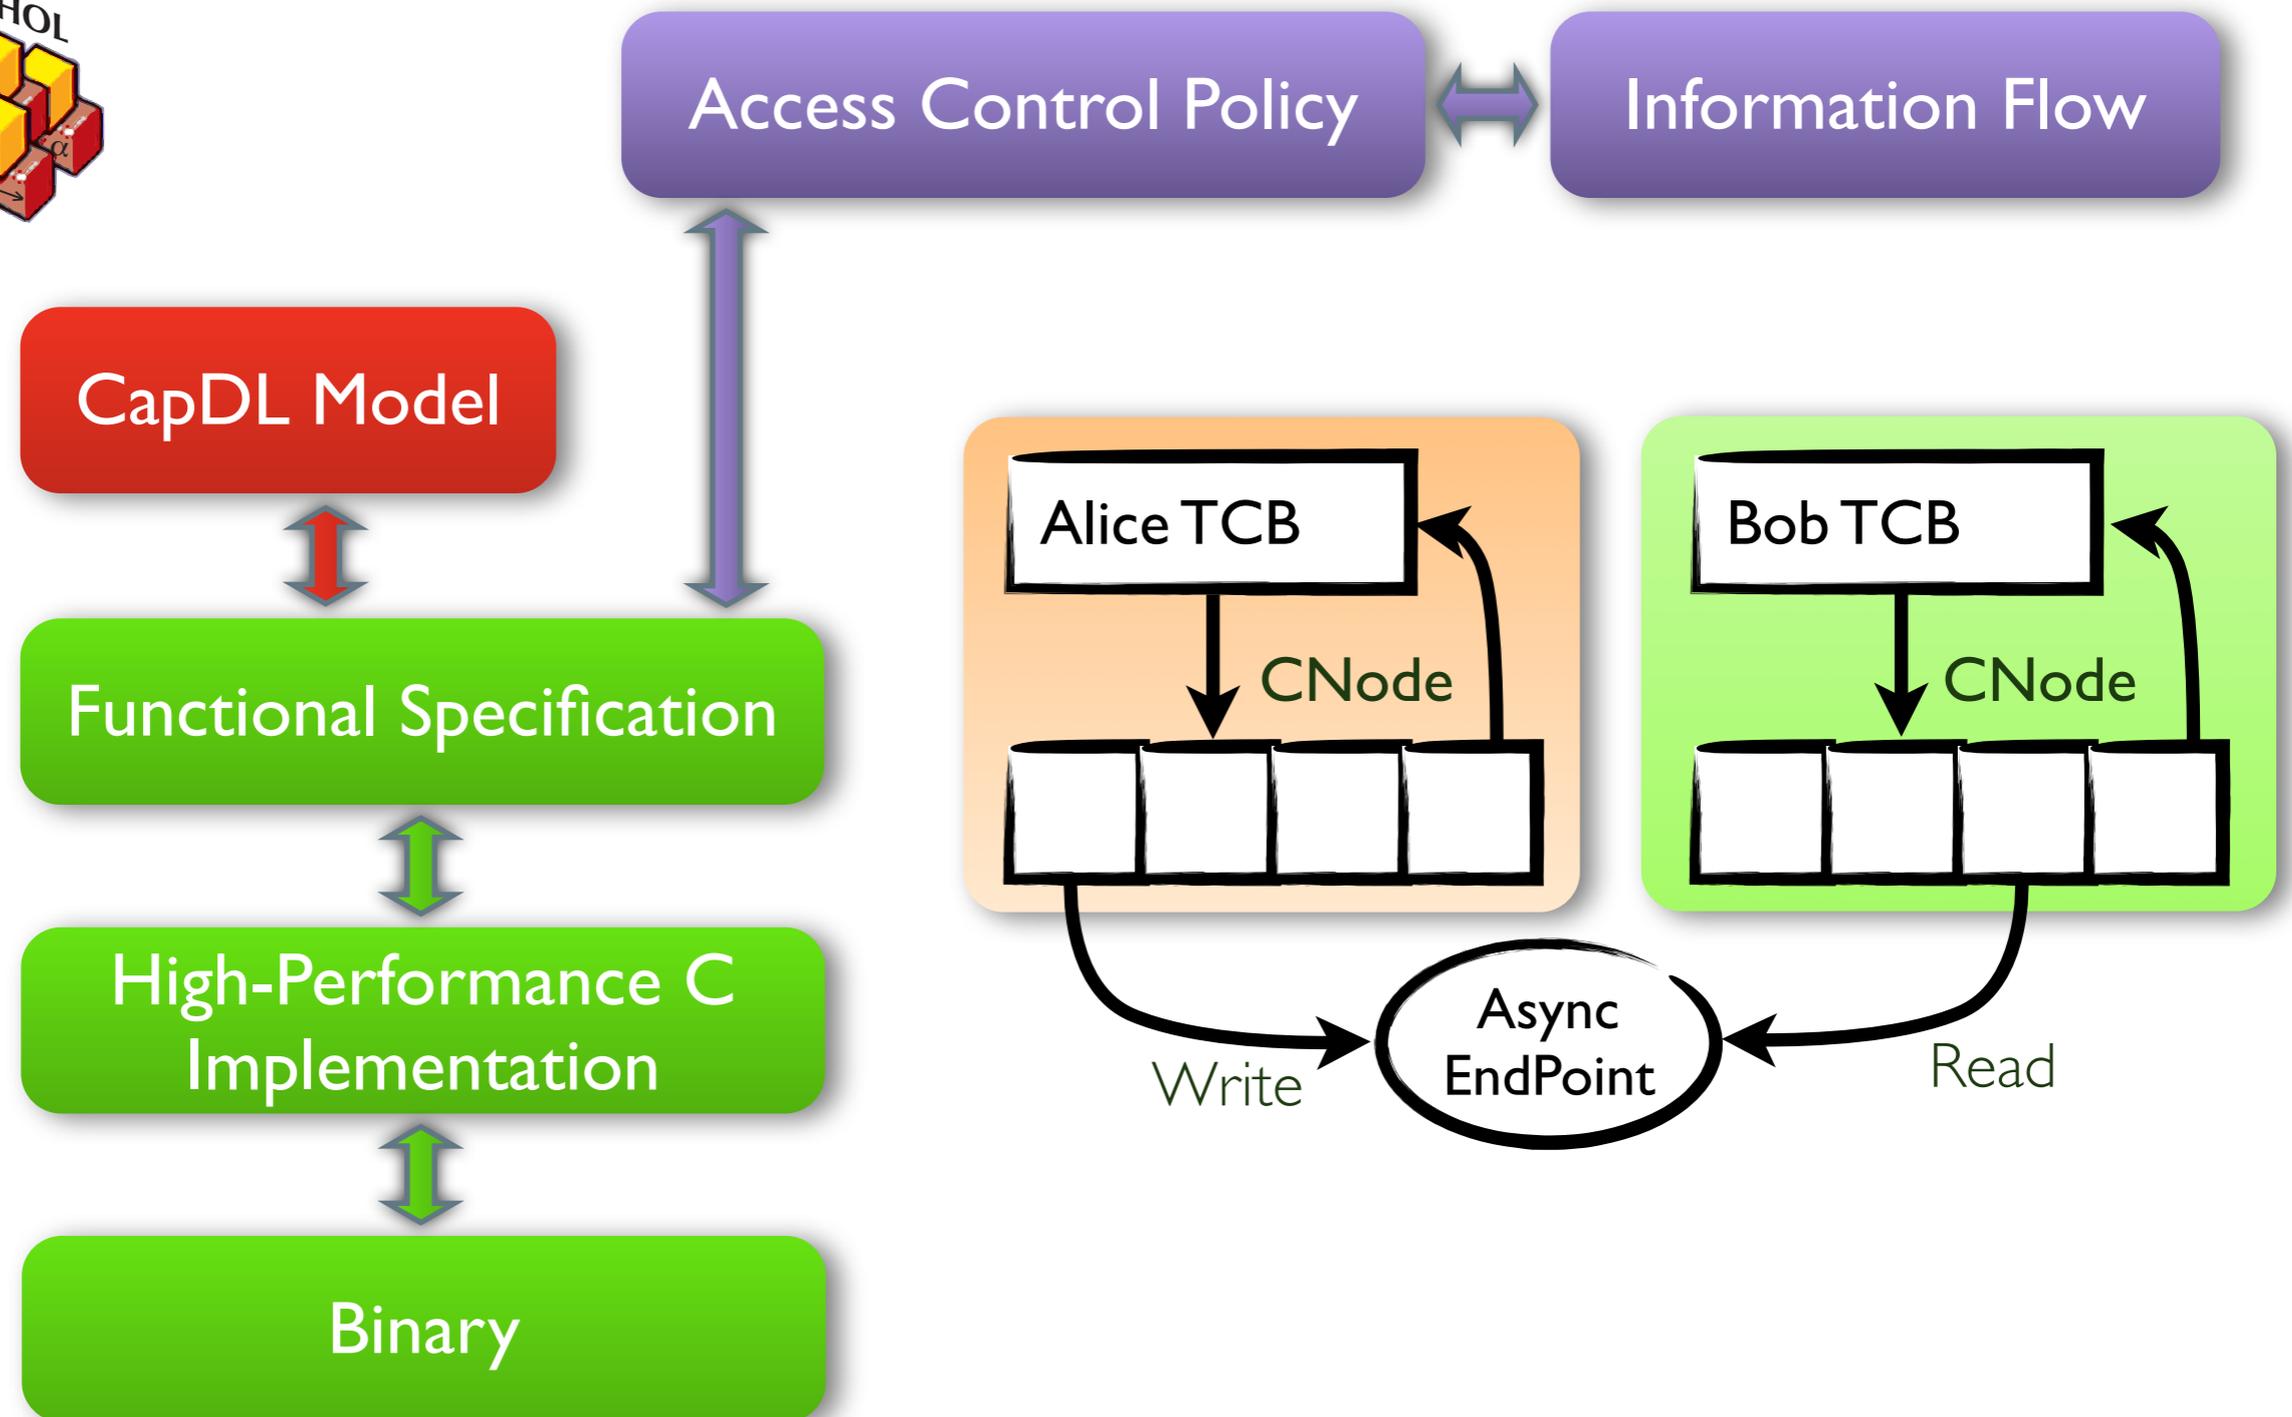
seL4 microkernel



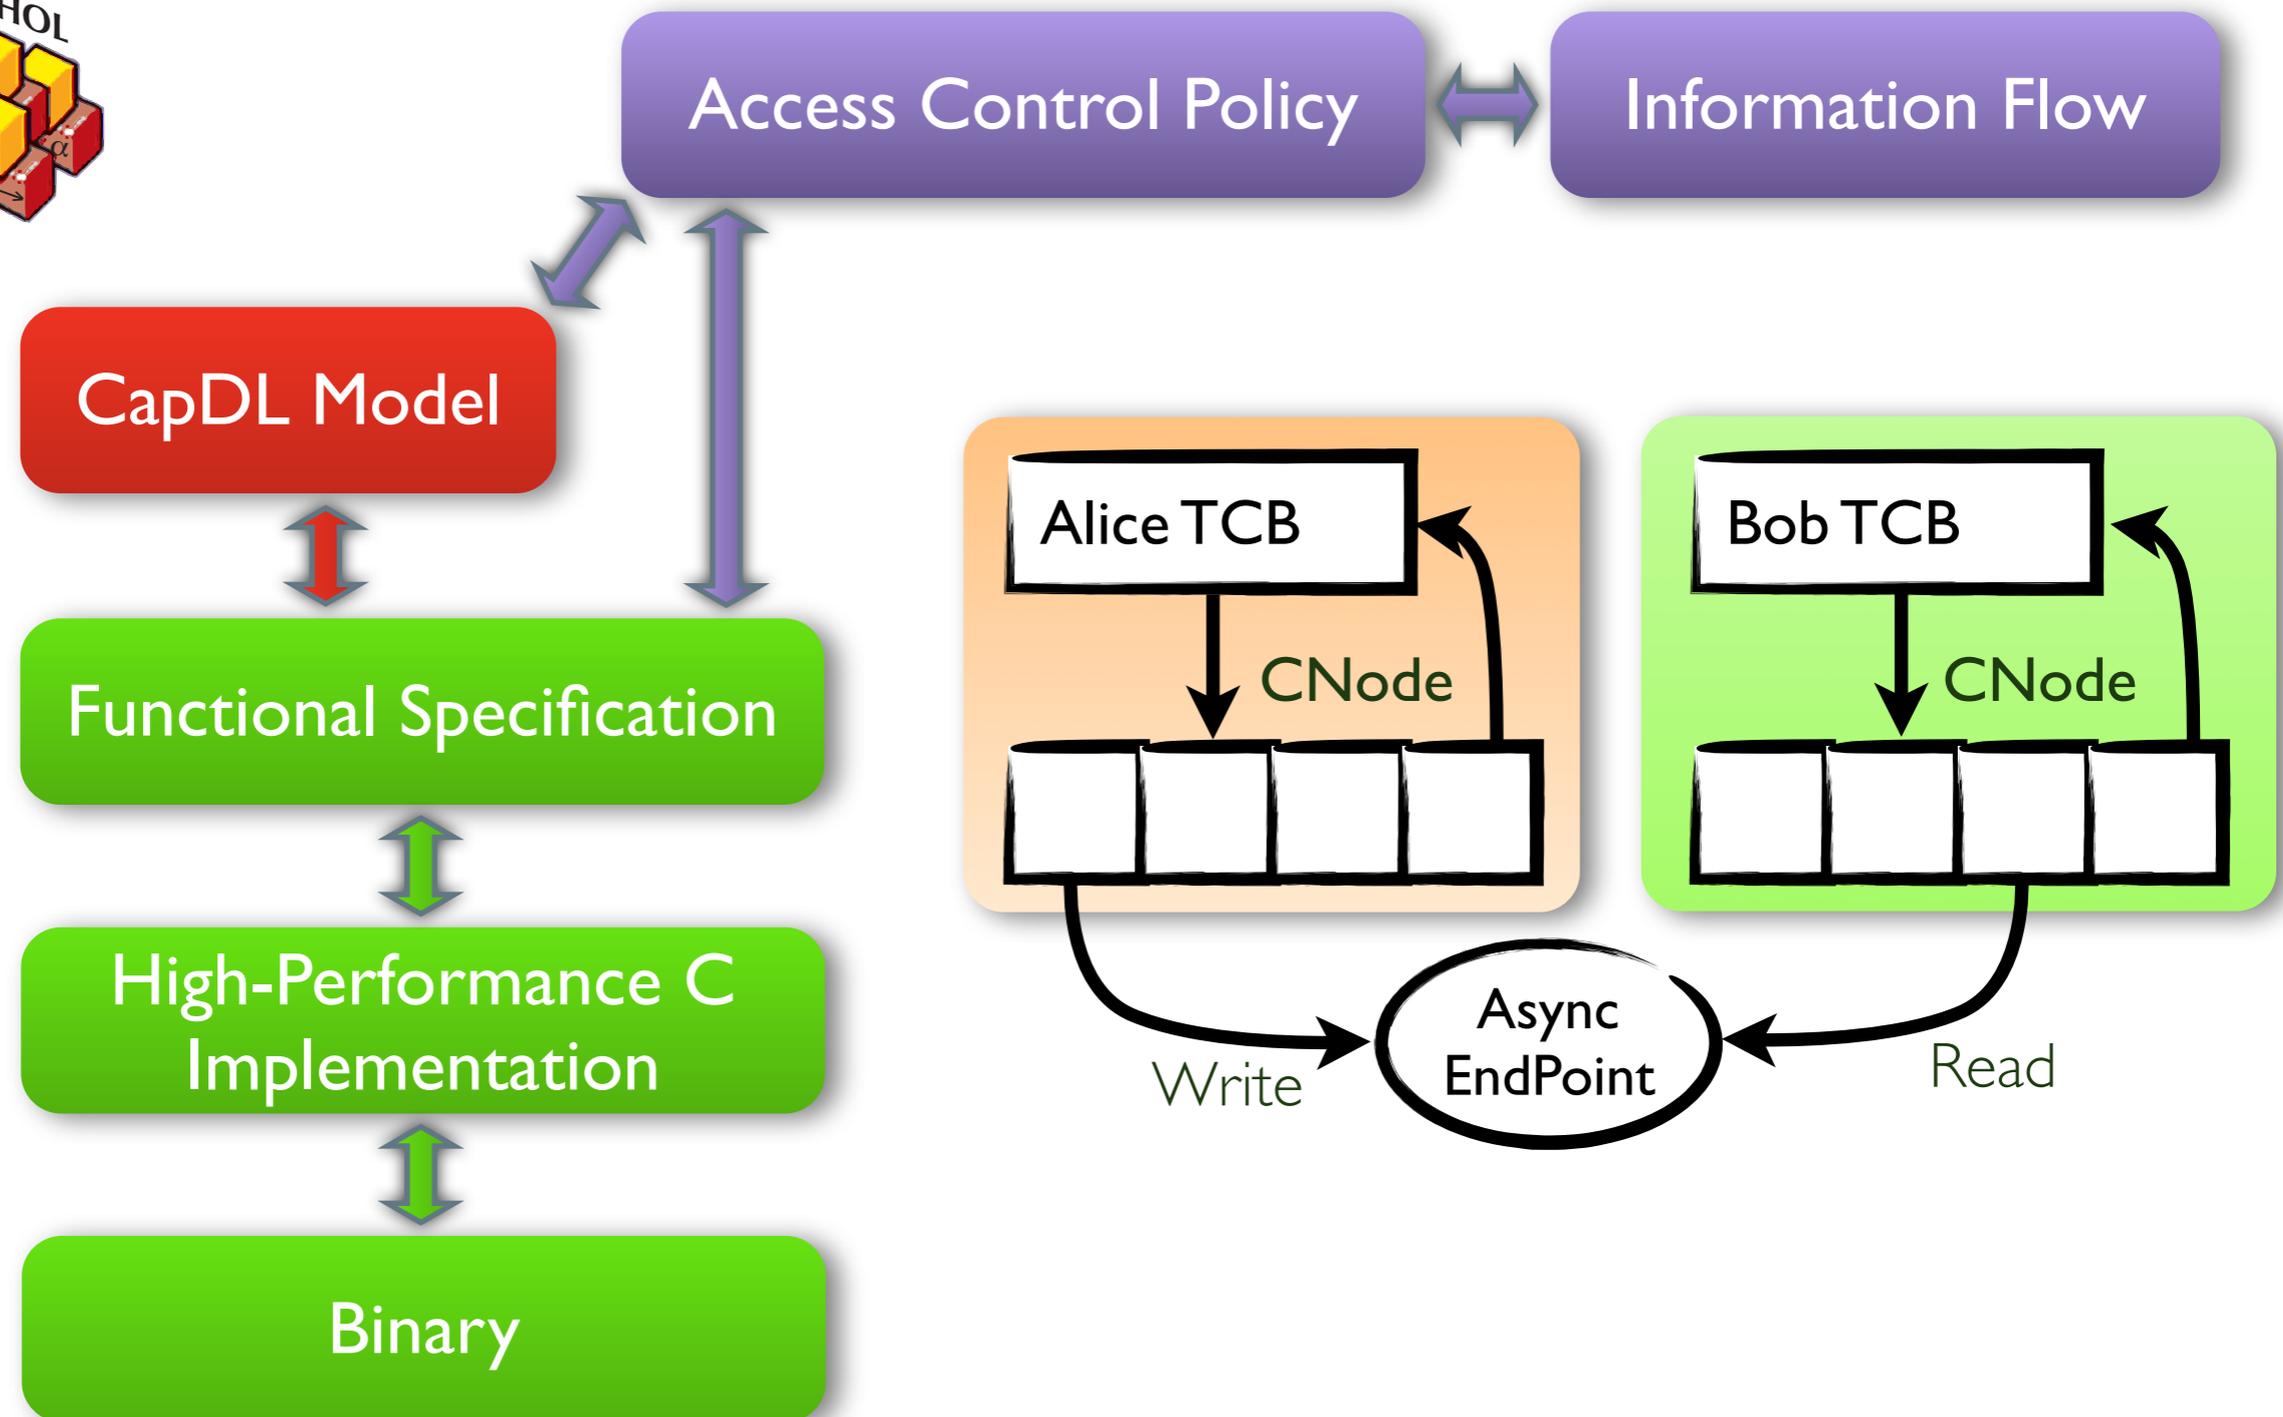
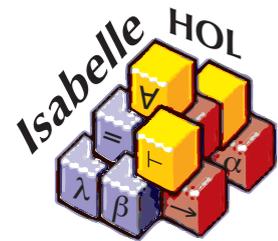
seL4 microkernel



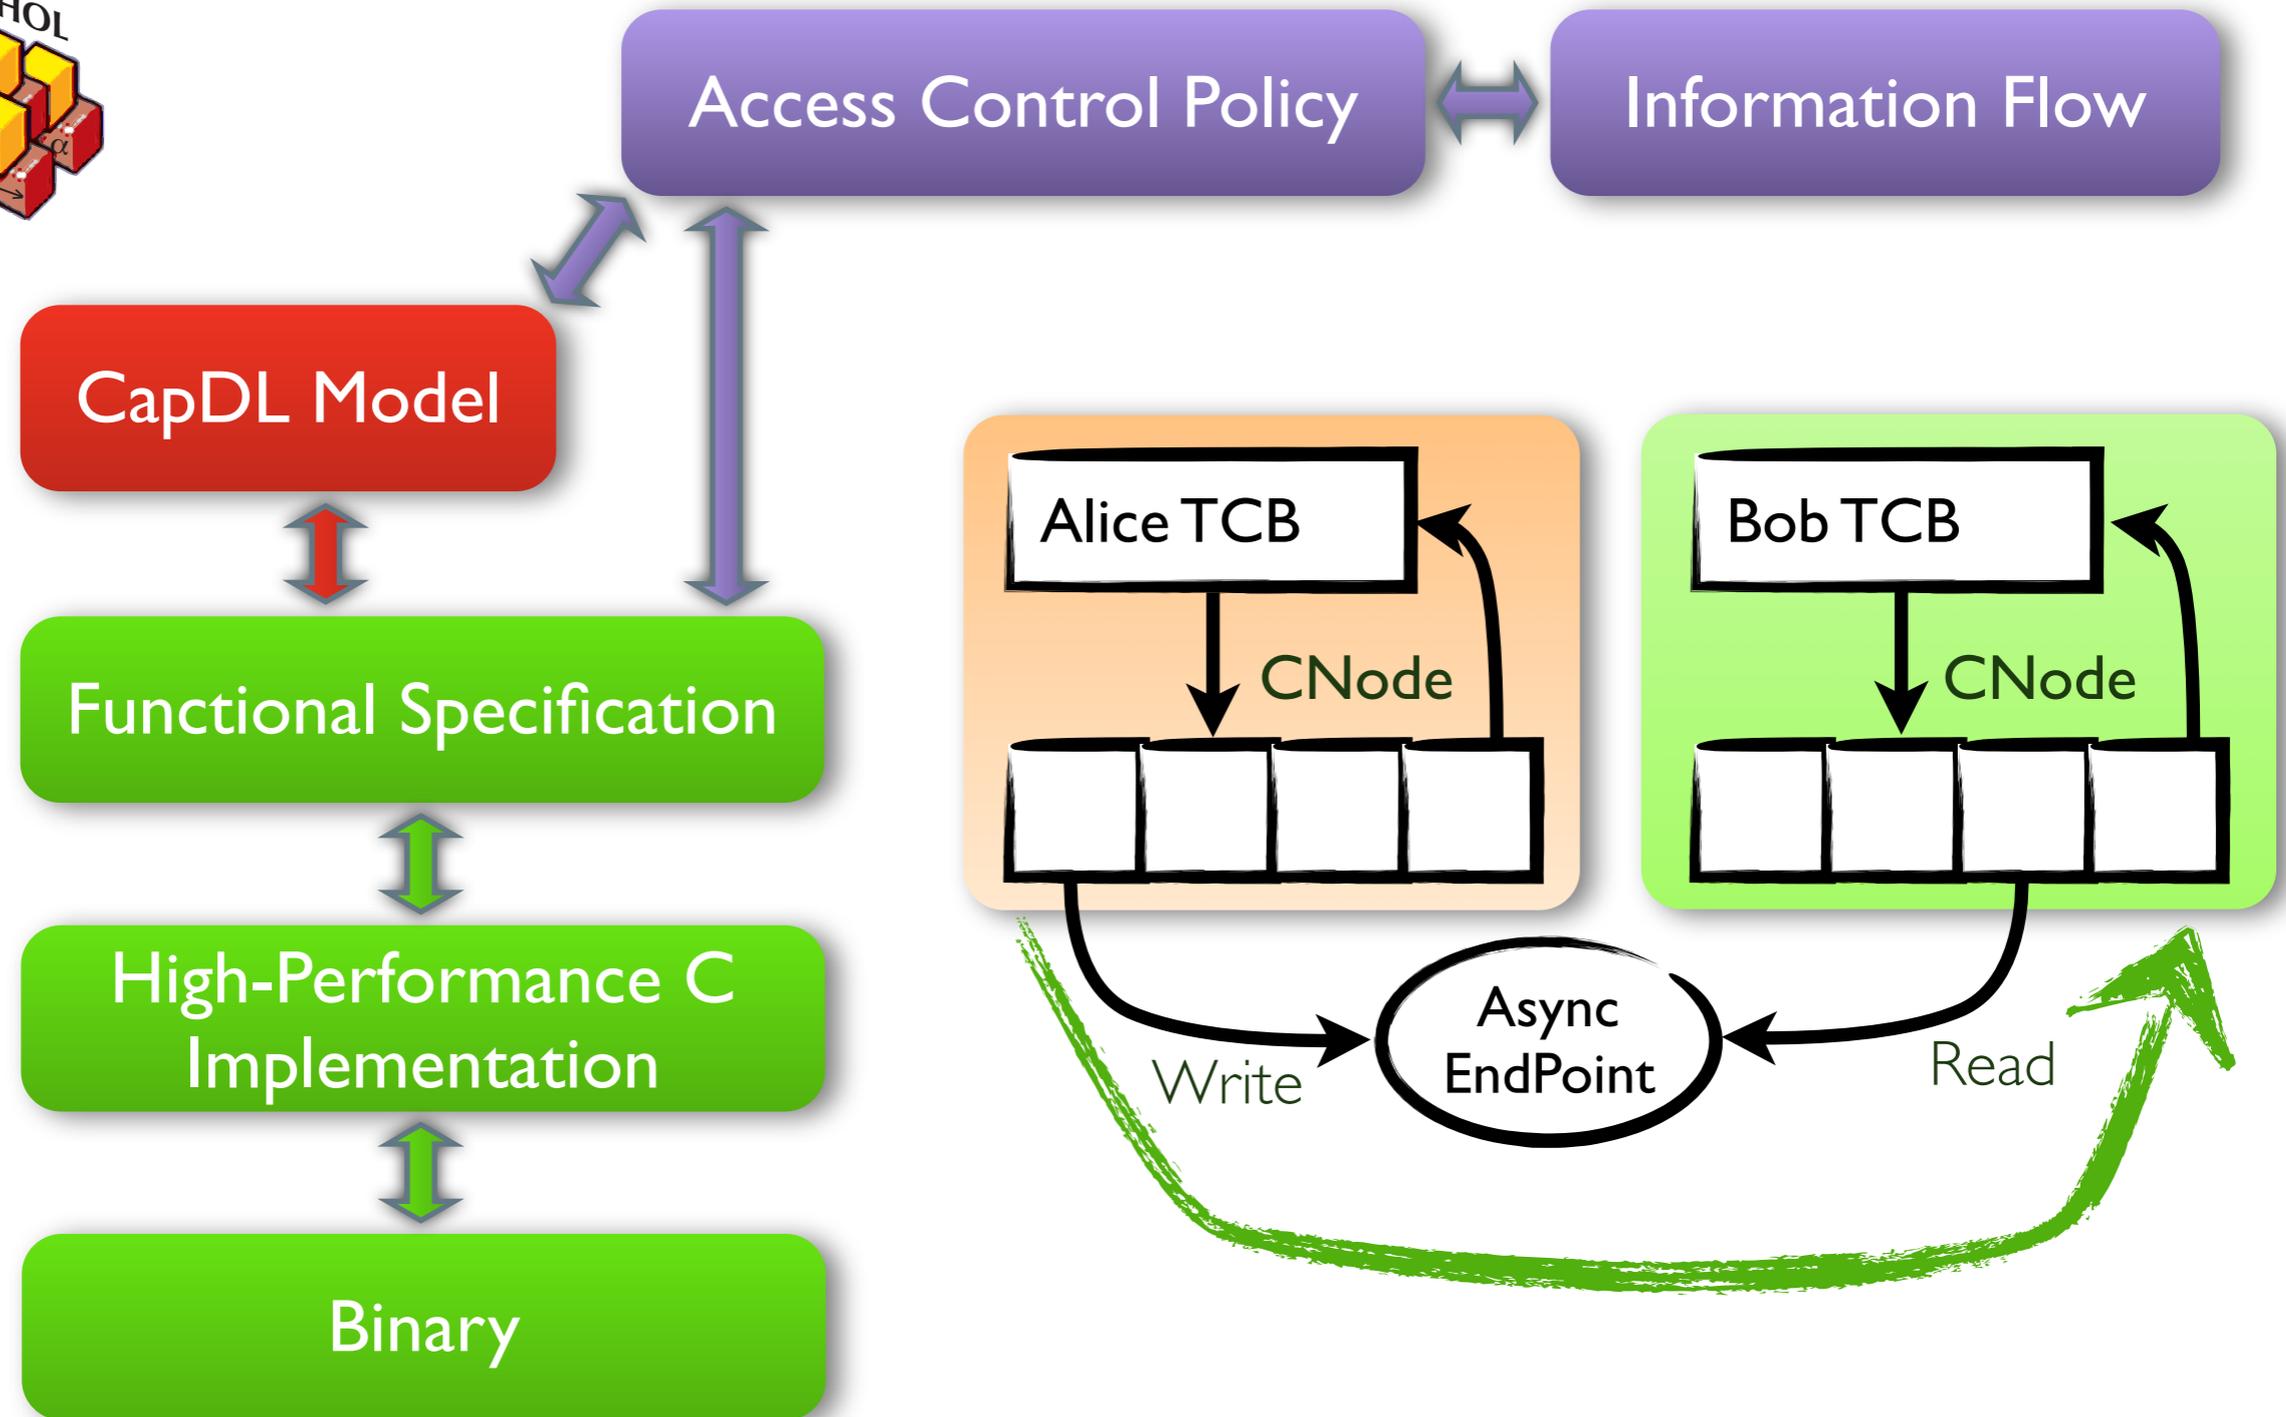
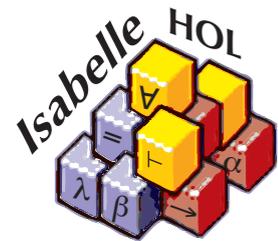
seL4 microkernel

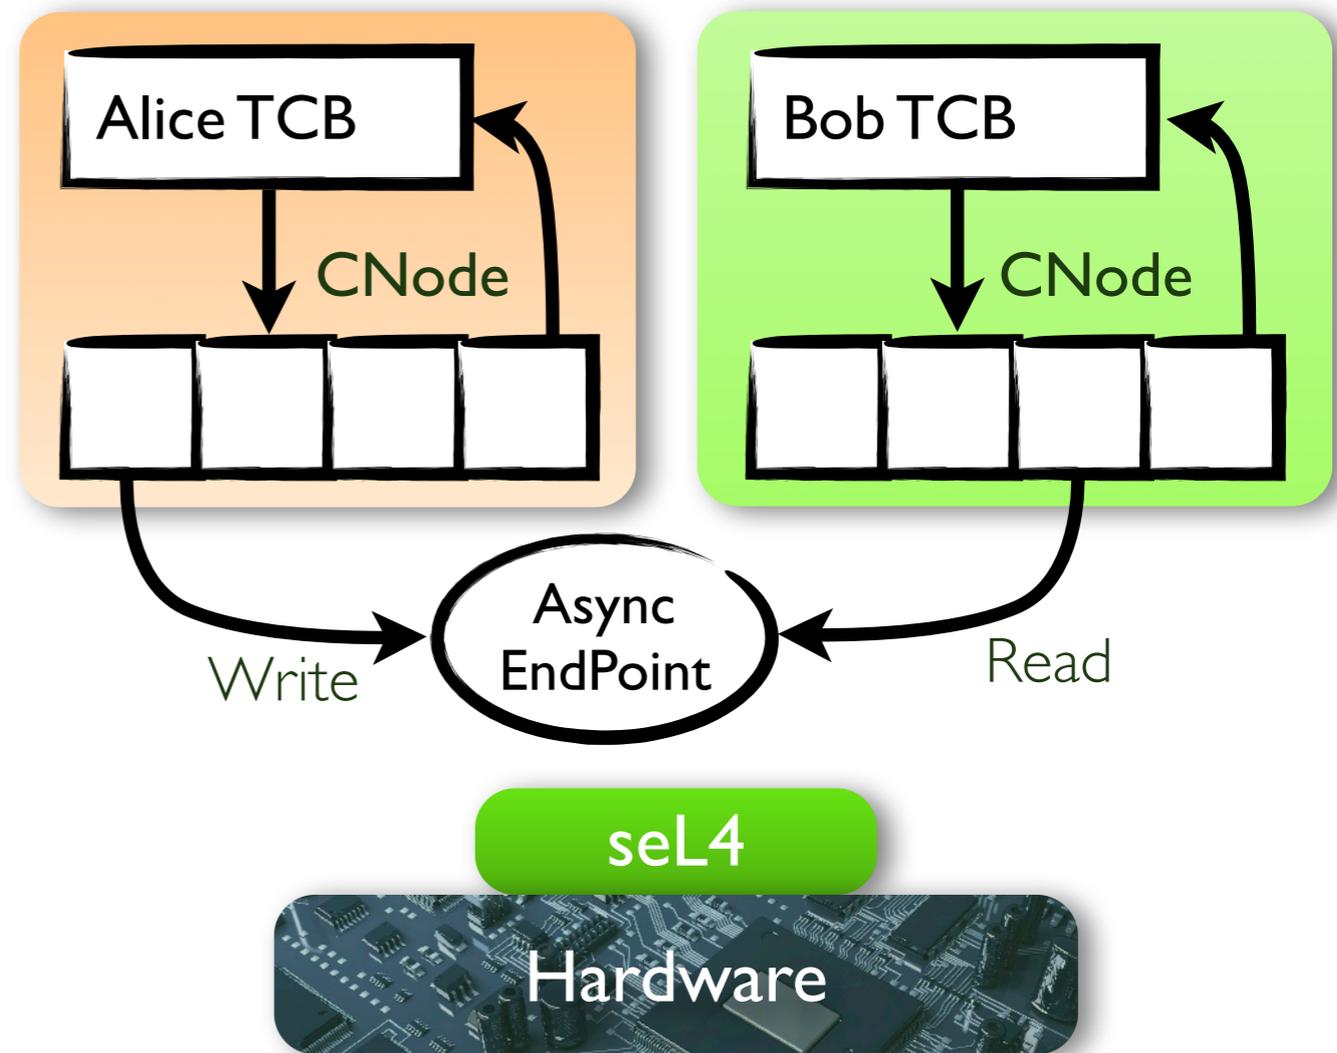


seL4 microkernel

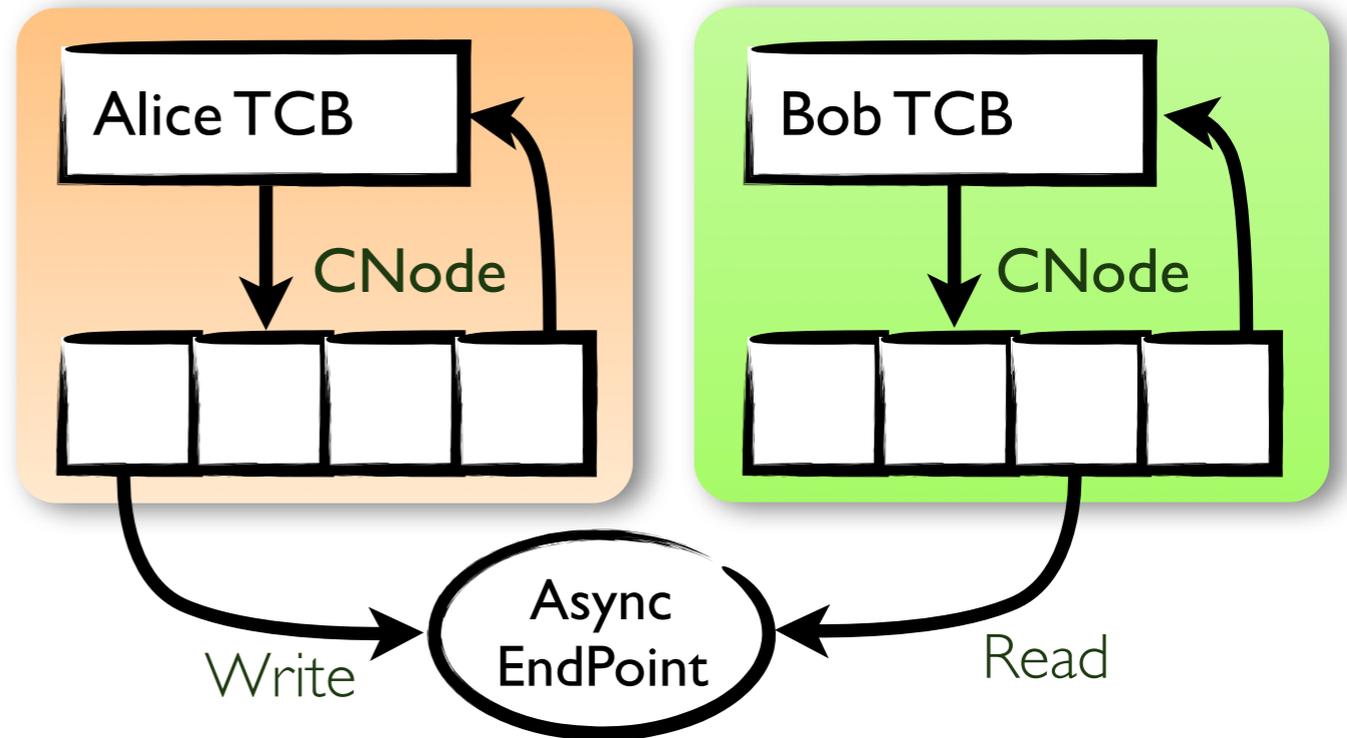
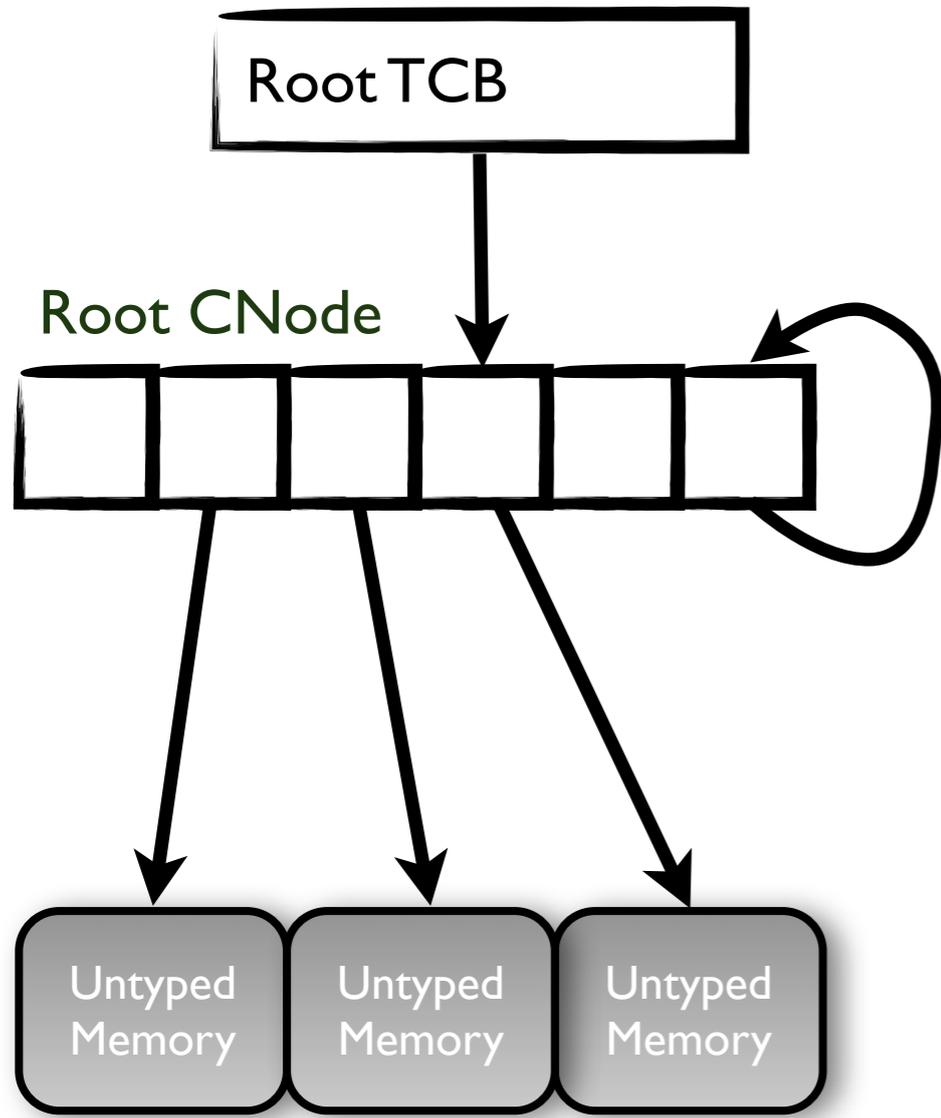


seL4 microkernel





Start State



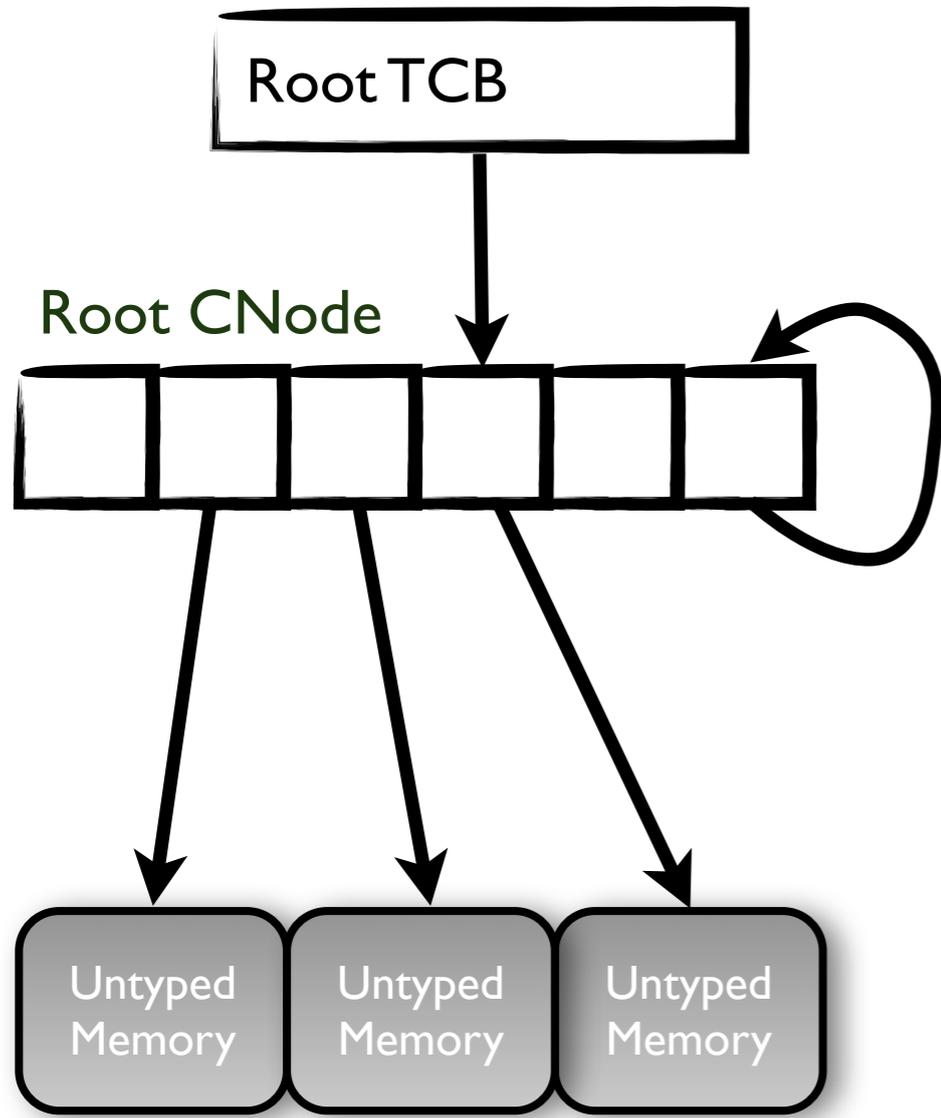
seL4

Hardware

seL4

Hardware

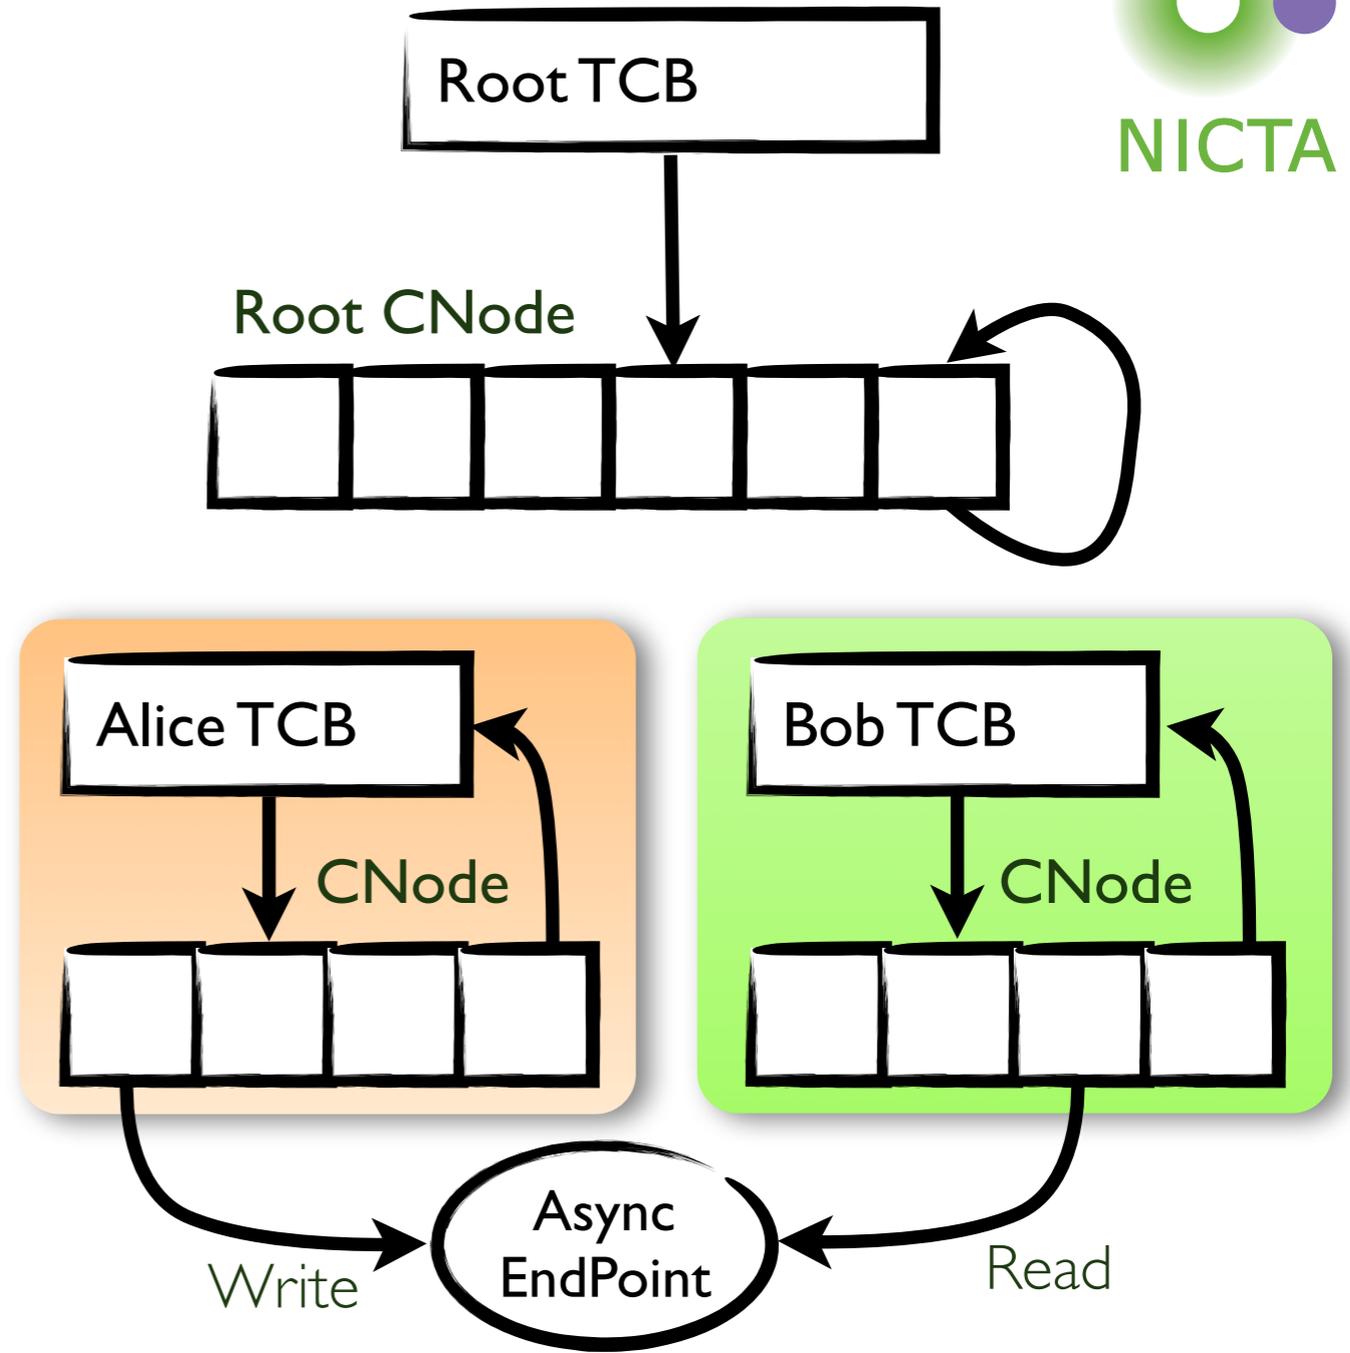
Start State



seL4



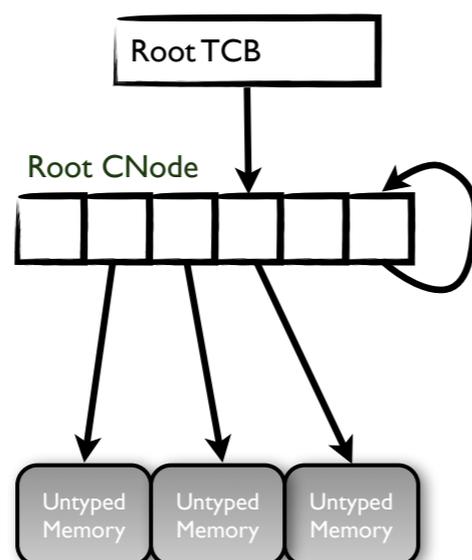
Initialised State



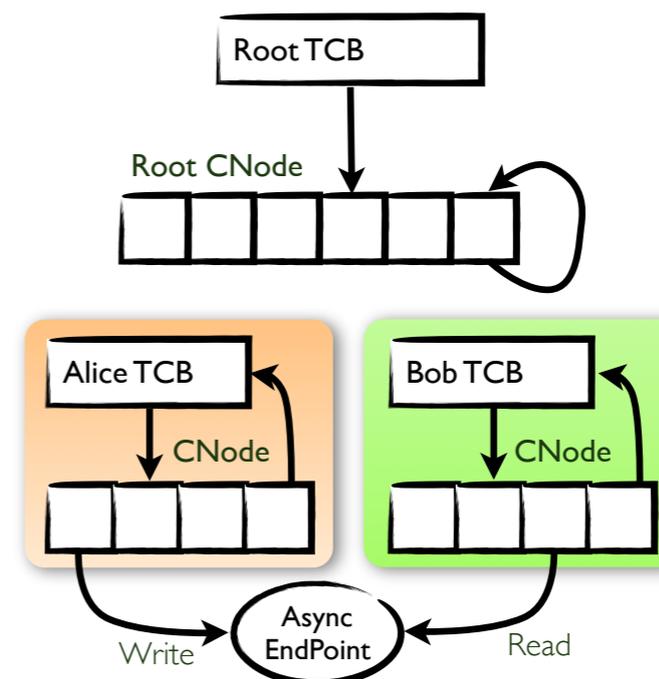
seL4

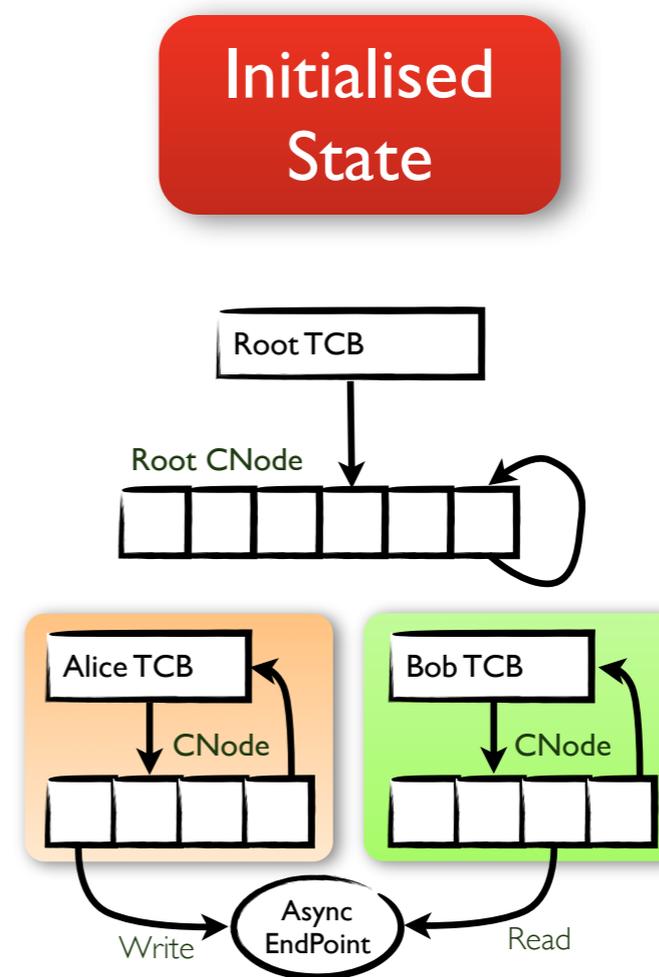
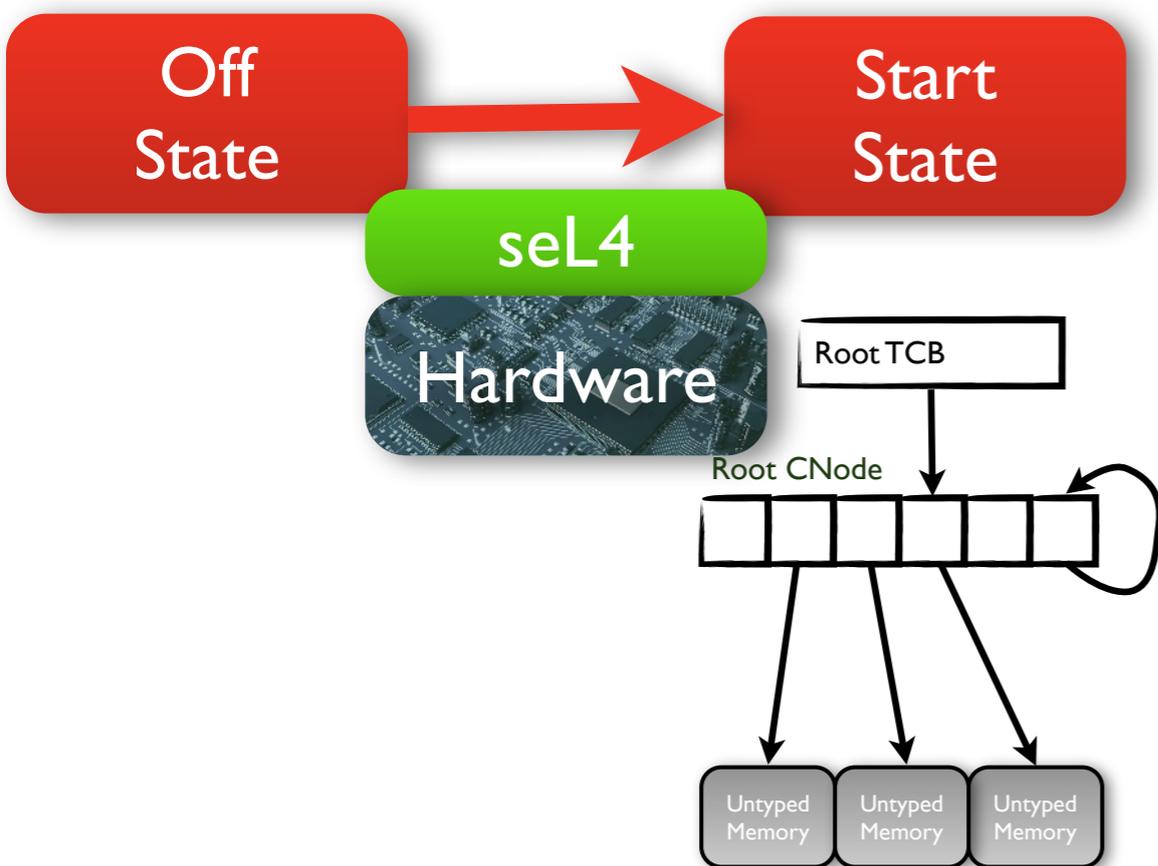


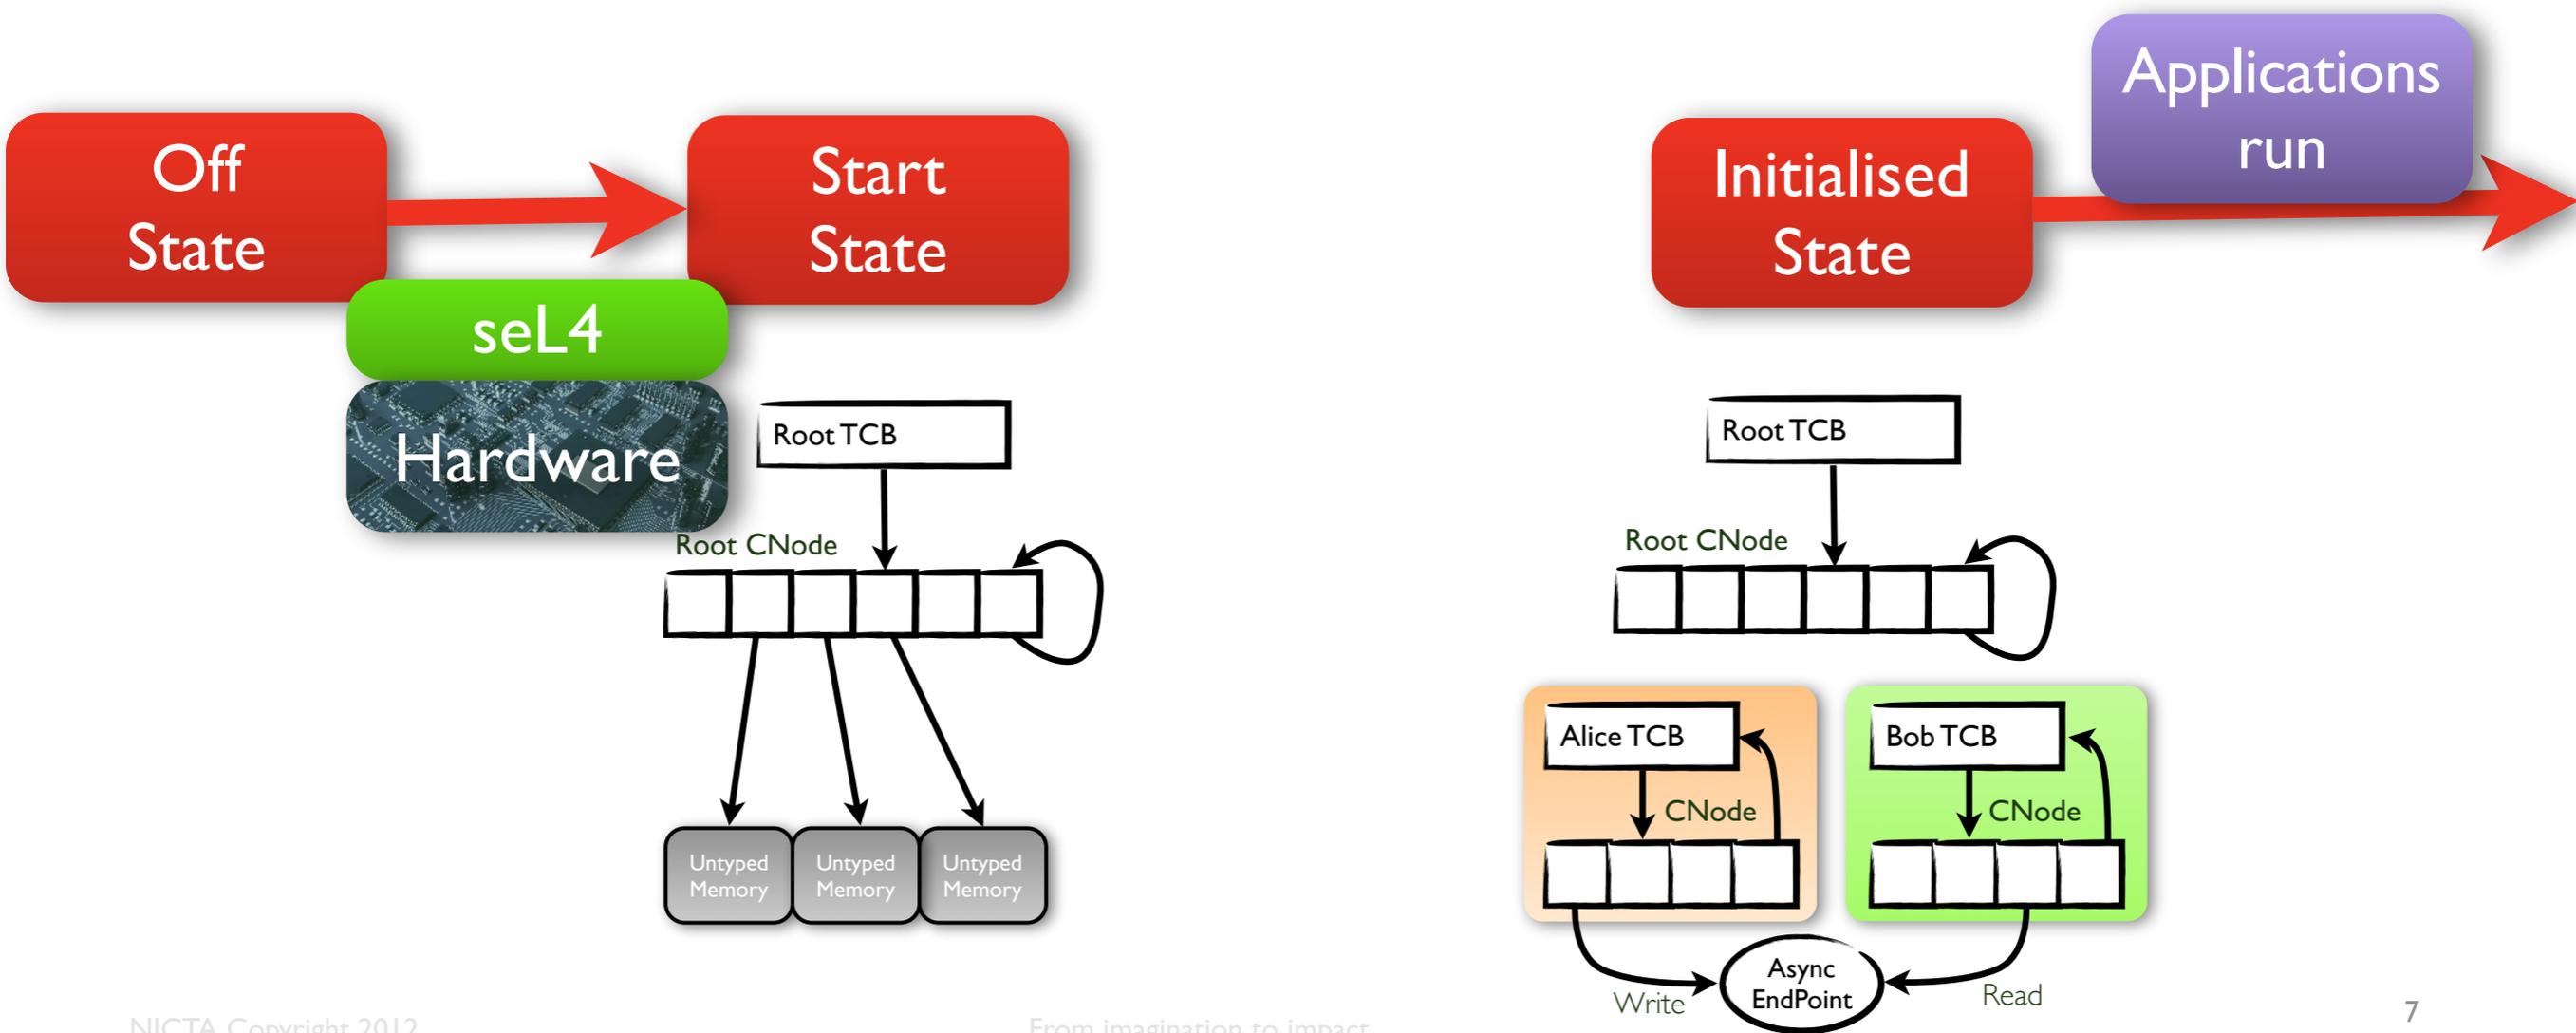
Start State

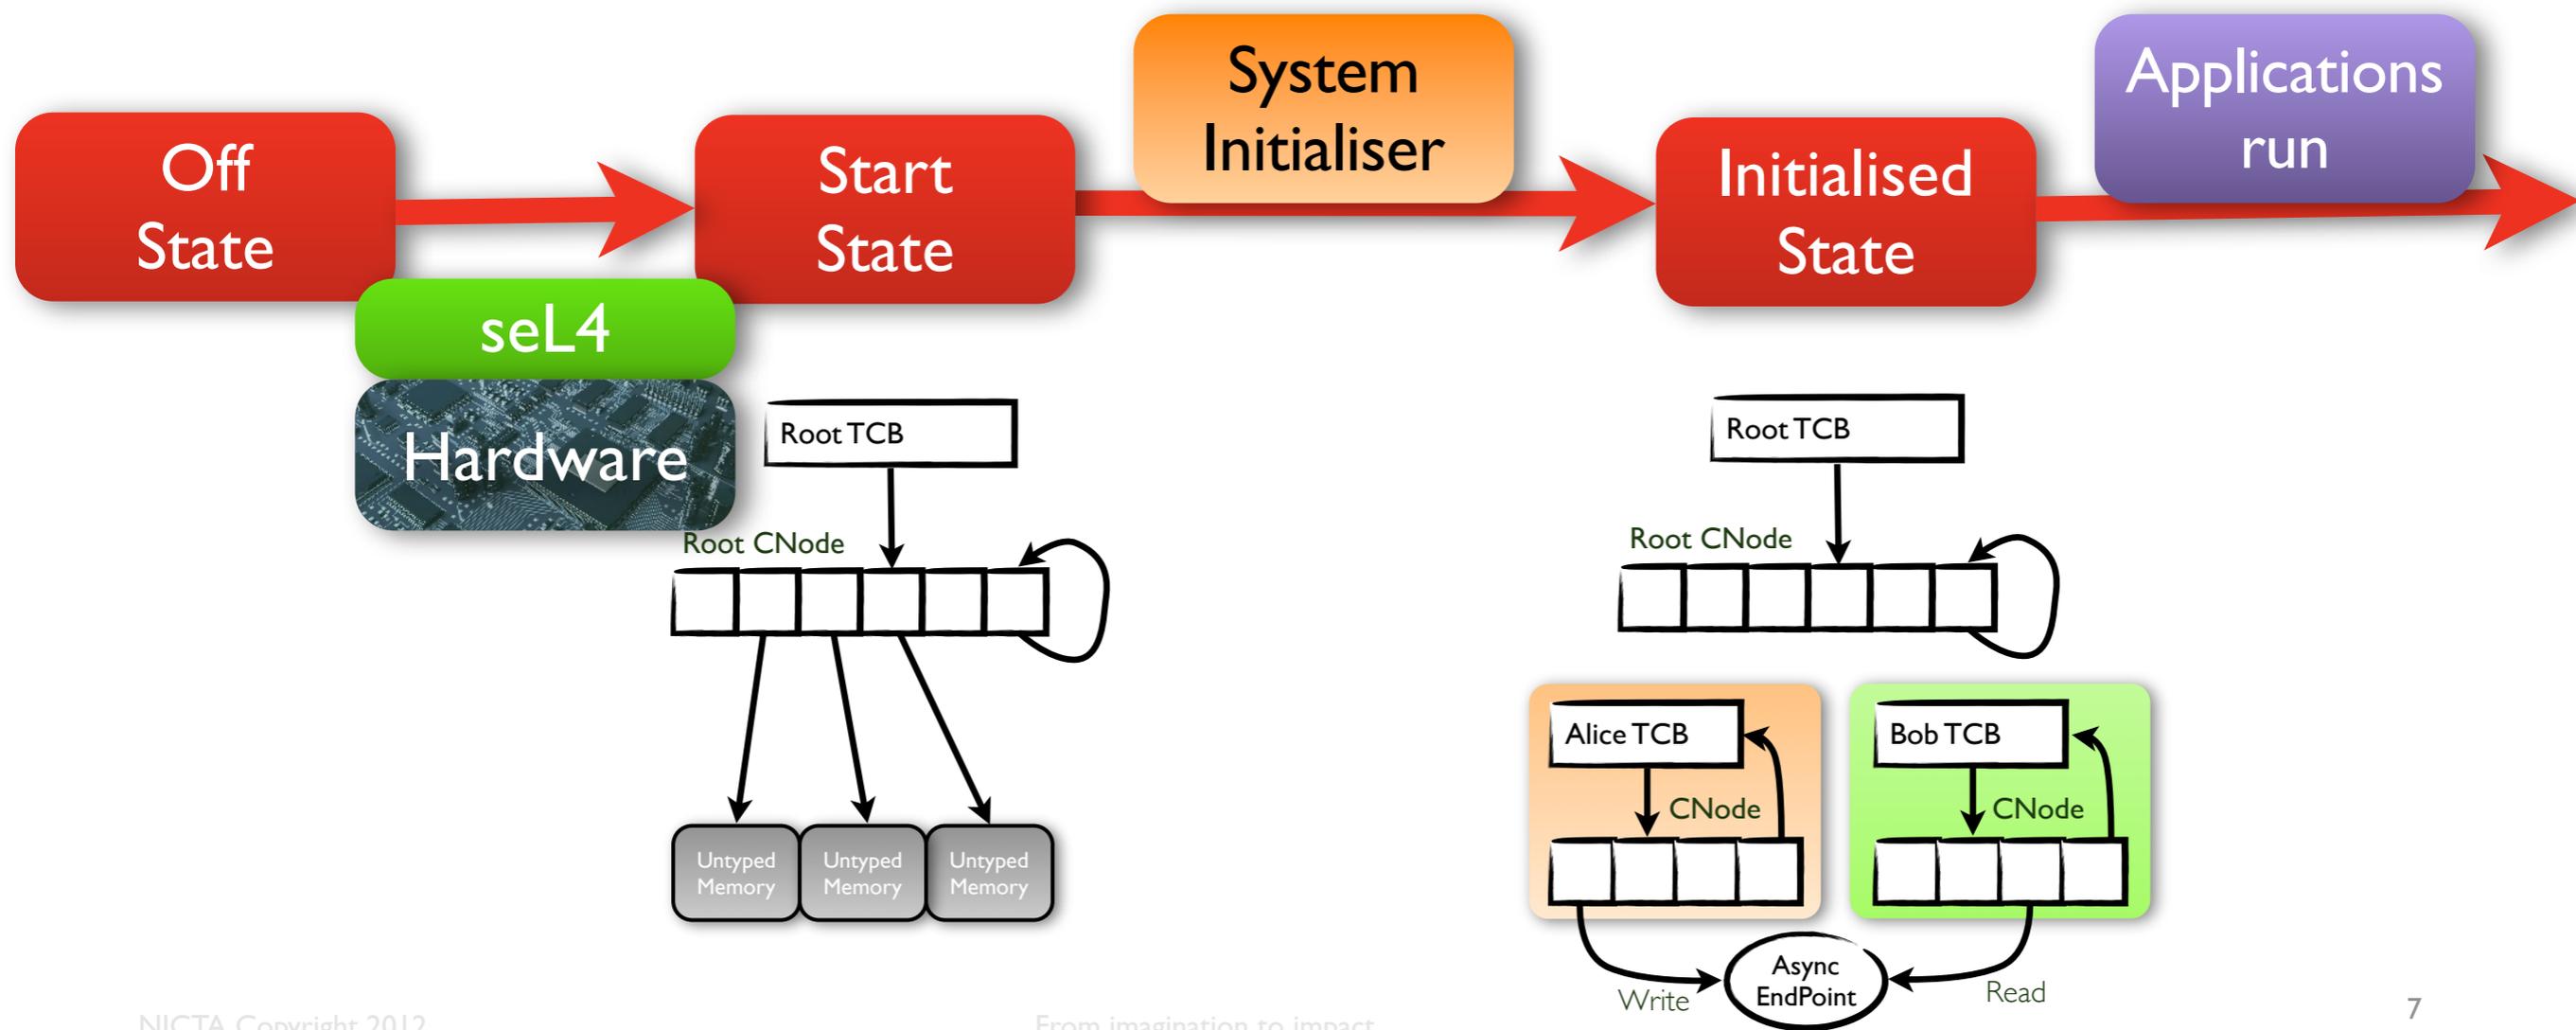


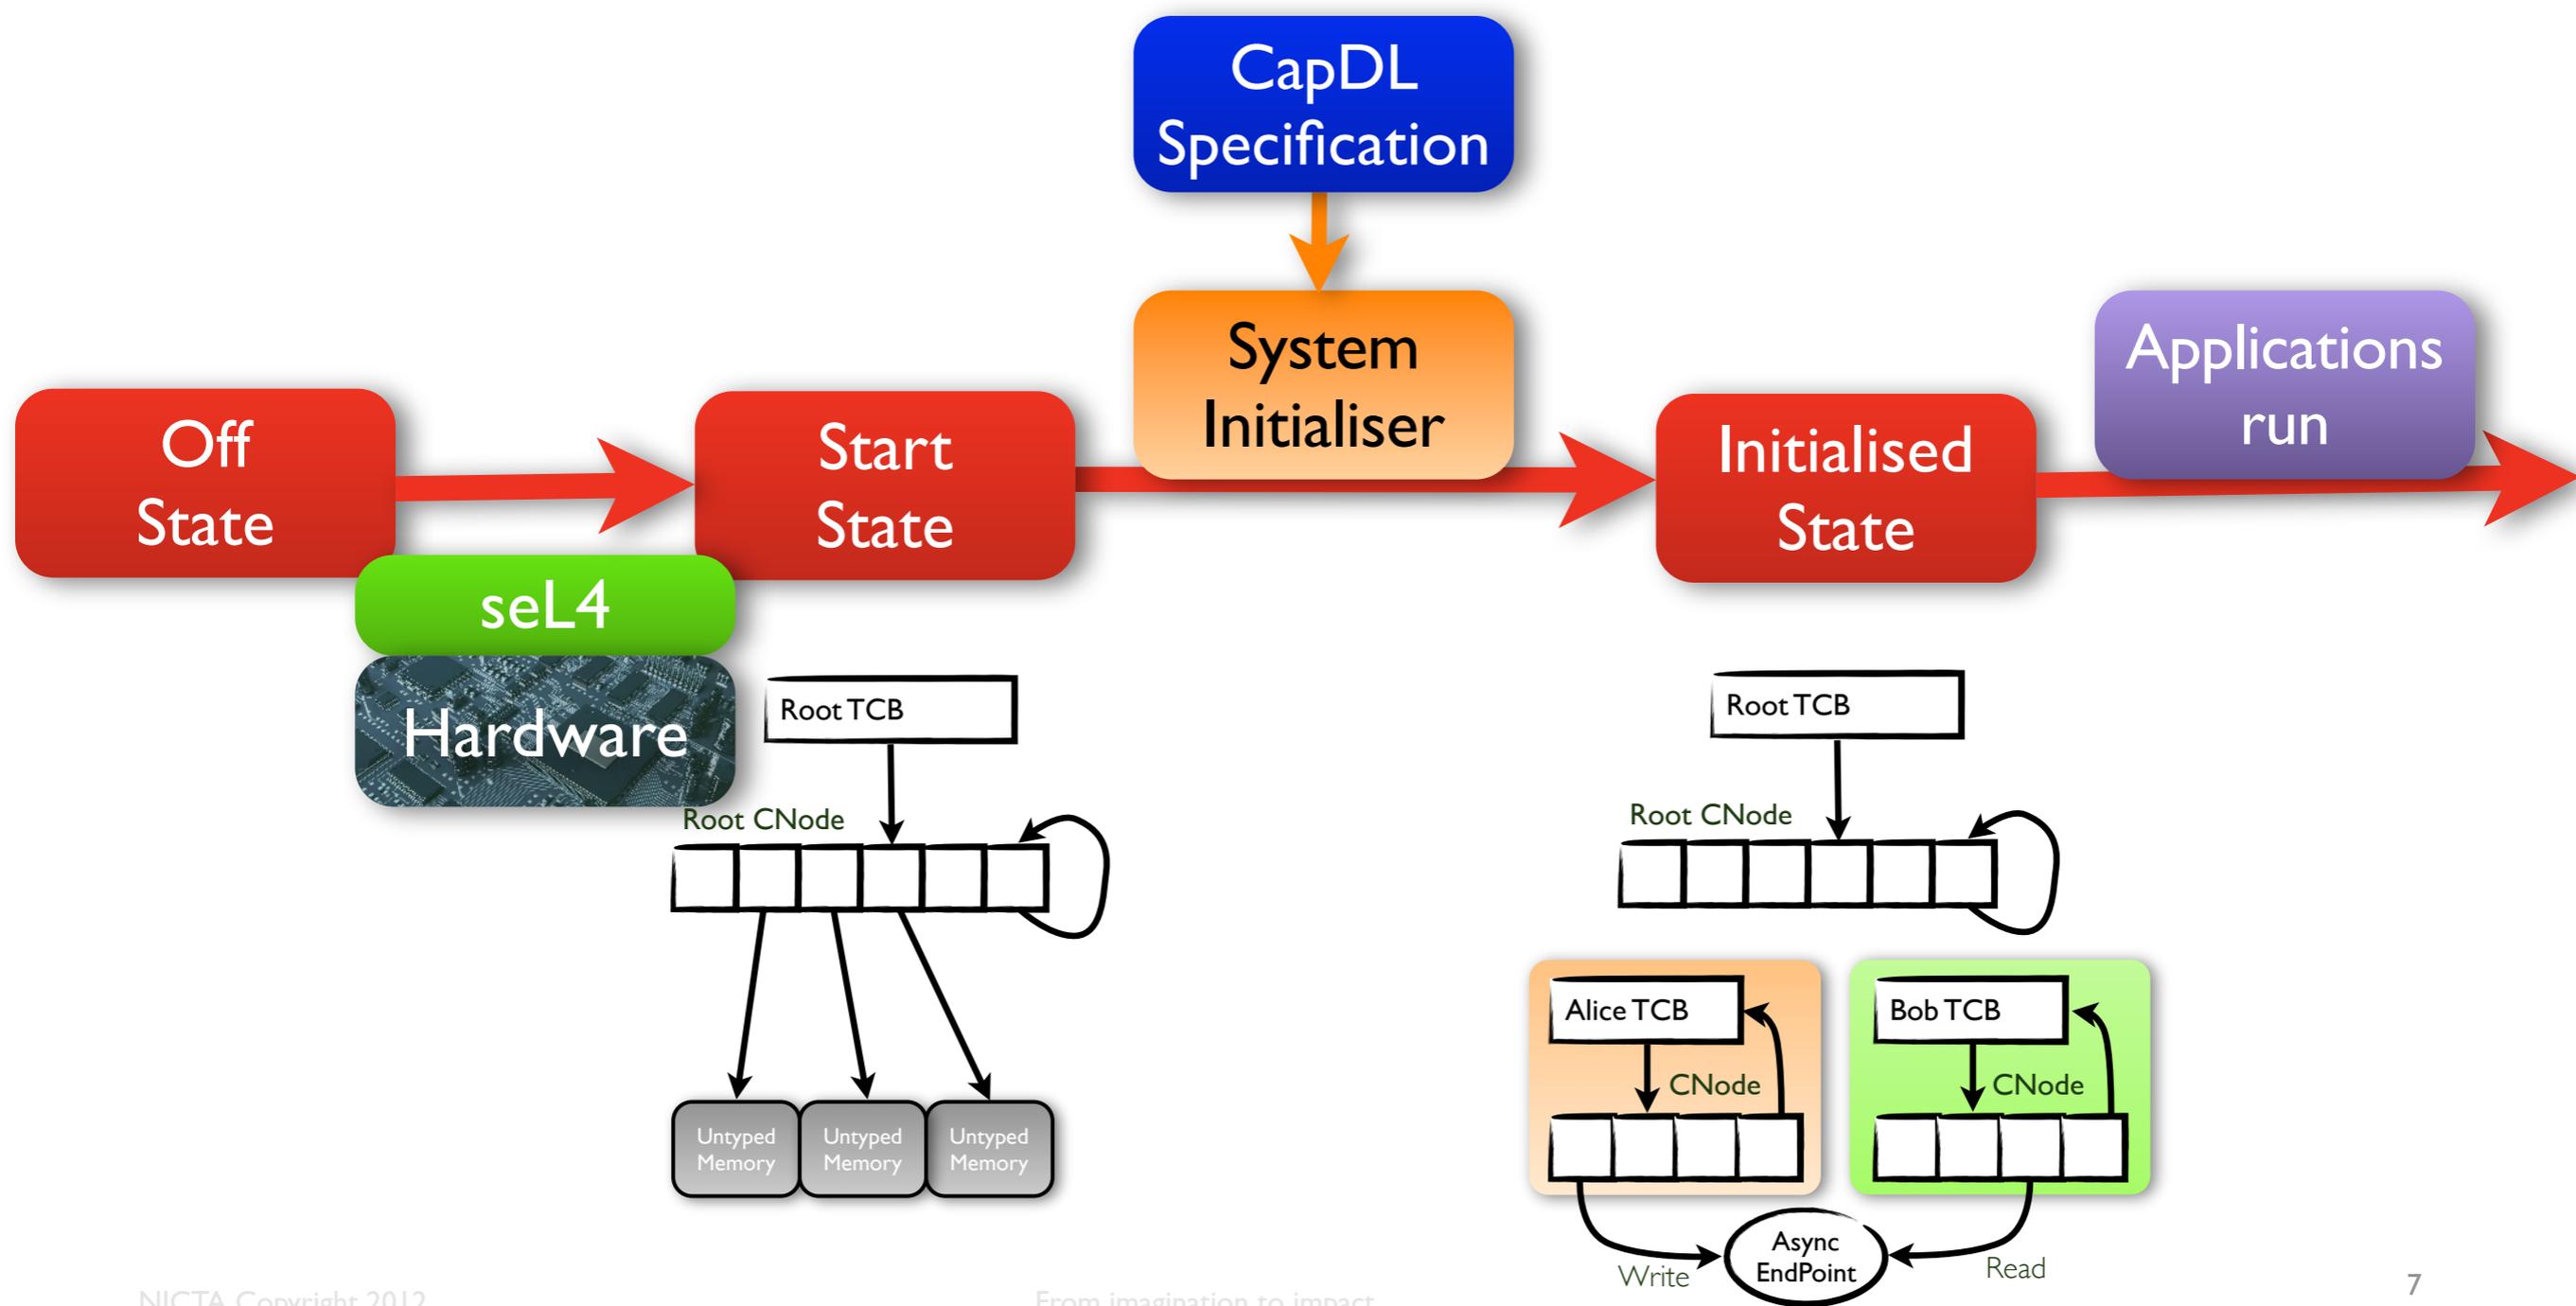
Initialised State











```

objects { tcb_alice = tcb ()
          cnode_alice = cnode (2 bits)
          cnode_bob = cnode (2 bits)
          tcb_bob = tcb ()
          aep_shared = aep }
caps { tcb_alice { cspace: cnode_alice (guard: 0, guard_size: 30) }
       tcb_bob { cspace: cnode_bob (guard: 0, guard_size: 30) }
       cnode_alice { 0x0: aep_shared (W) }
       cnode_bob { 0x2: aep_shared (R) } }
  
```

CapDL Specification

System Initialiser

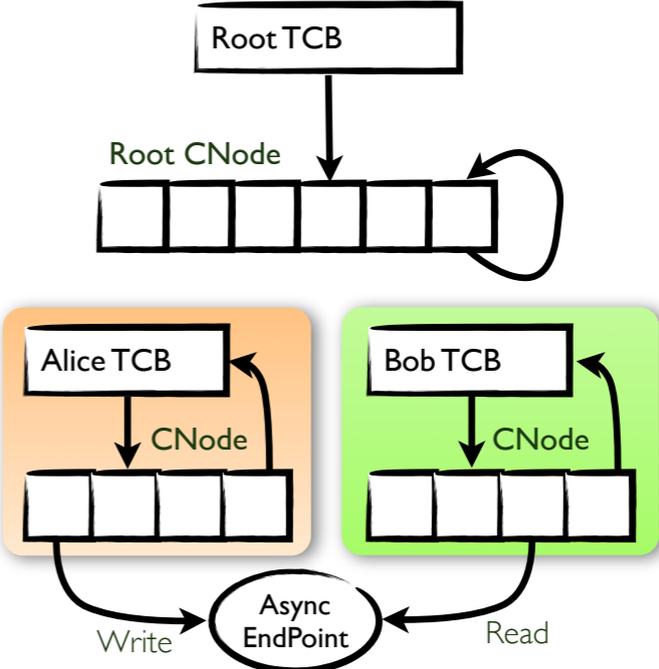
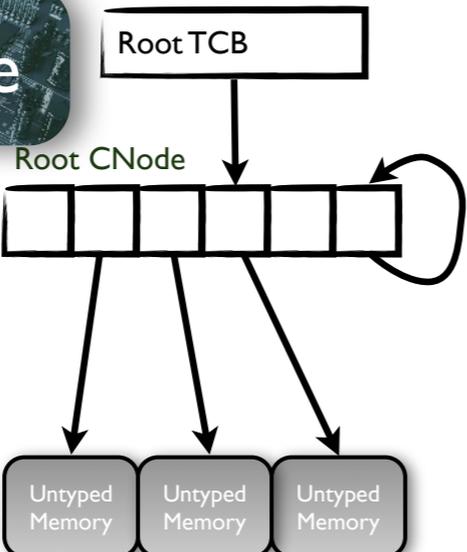
Applications run

Off State

Start State

Initialised State

seL4



```

objects { tcb_alice = tcb ()
          cnode_alice = cnode (2 bits)
          cnode_bob = cnode (2 bits)
          tcb_bob = tcb ()
          aep_shared = aep }
caps { tcb_alice { cspace: cnode_alice (guard: 0, guard_size: 30) }
       tcb_bob { cspace: cnode_bob (guard: 0, guard_size: 30) }
       cnode_alice { 0x0: aep_shared (W) }
       cnode_bob { 0x2: aep_shared (R) } }
  
```

CapDL Specification

System Initialiser

Equivalent

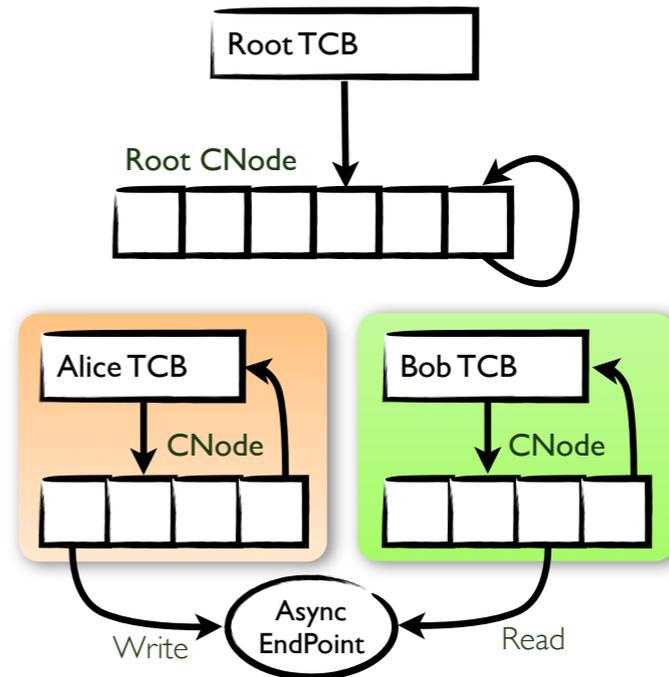
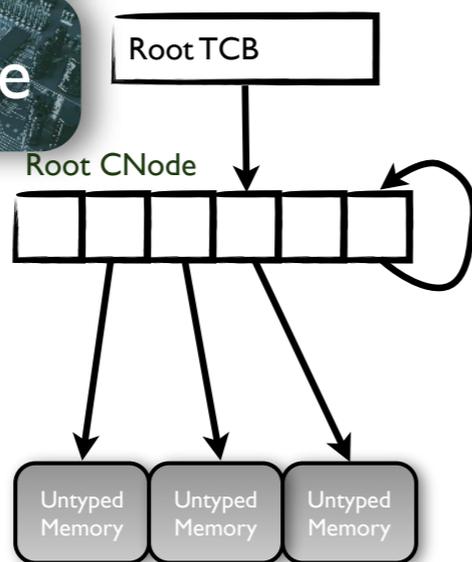
Off State

Start State

Initialised State

Applications run

seL4



```

objects { tcb_alice = tcb ()
          cnode_alice = cnode (2 bits)
          cnode_bob = cnode (2 bits)
          tcb_bob = tcb ()
          aep_shared = aep }
caps { tcb_alice { cspace: cnode_alice (guard: 0, guard_size: 30) }
       tcb_bob { cspace: cnode_bob (guard: 0, guard_size: 30) }
       cnode_alice { 0x0: aep_shared (W) }
       cnode_bob { 0x2: aep_shared (R) } }

```

```

{ valid_boot_info boot_info spec }
init_system spec boot_info
{ objects_initialised spec }

```

CapDL Specification

Equivalent

System Initialiser

Applications run

Off State

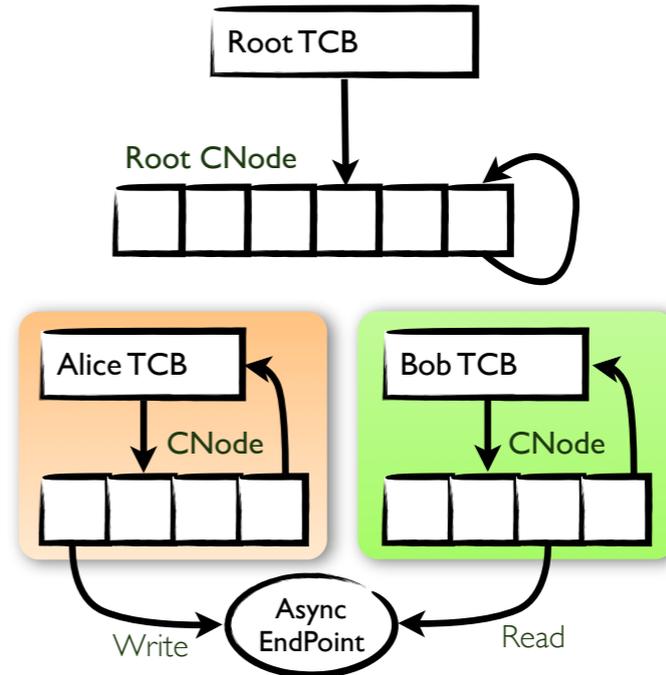
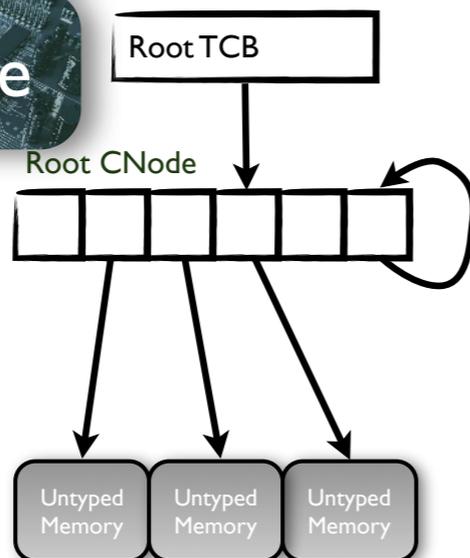
Start State

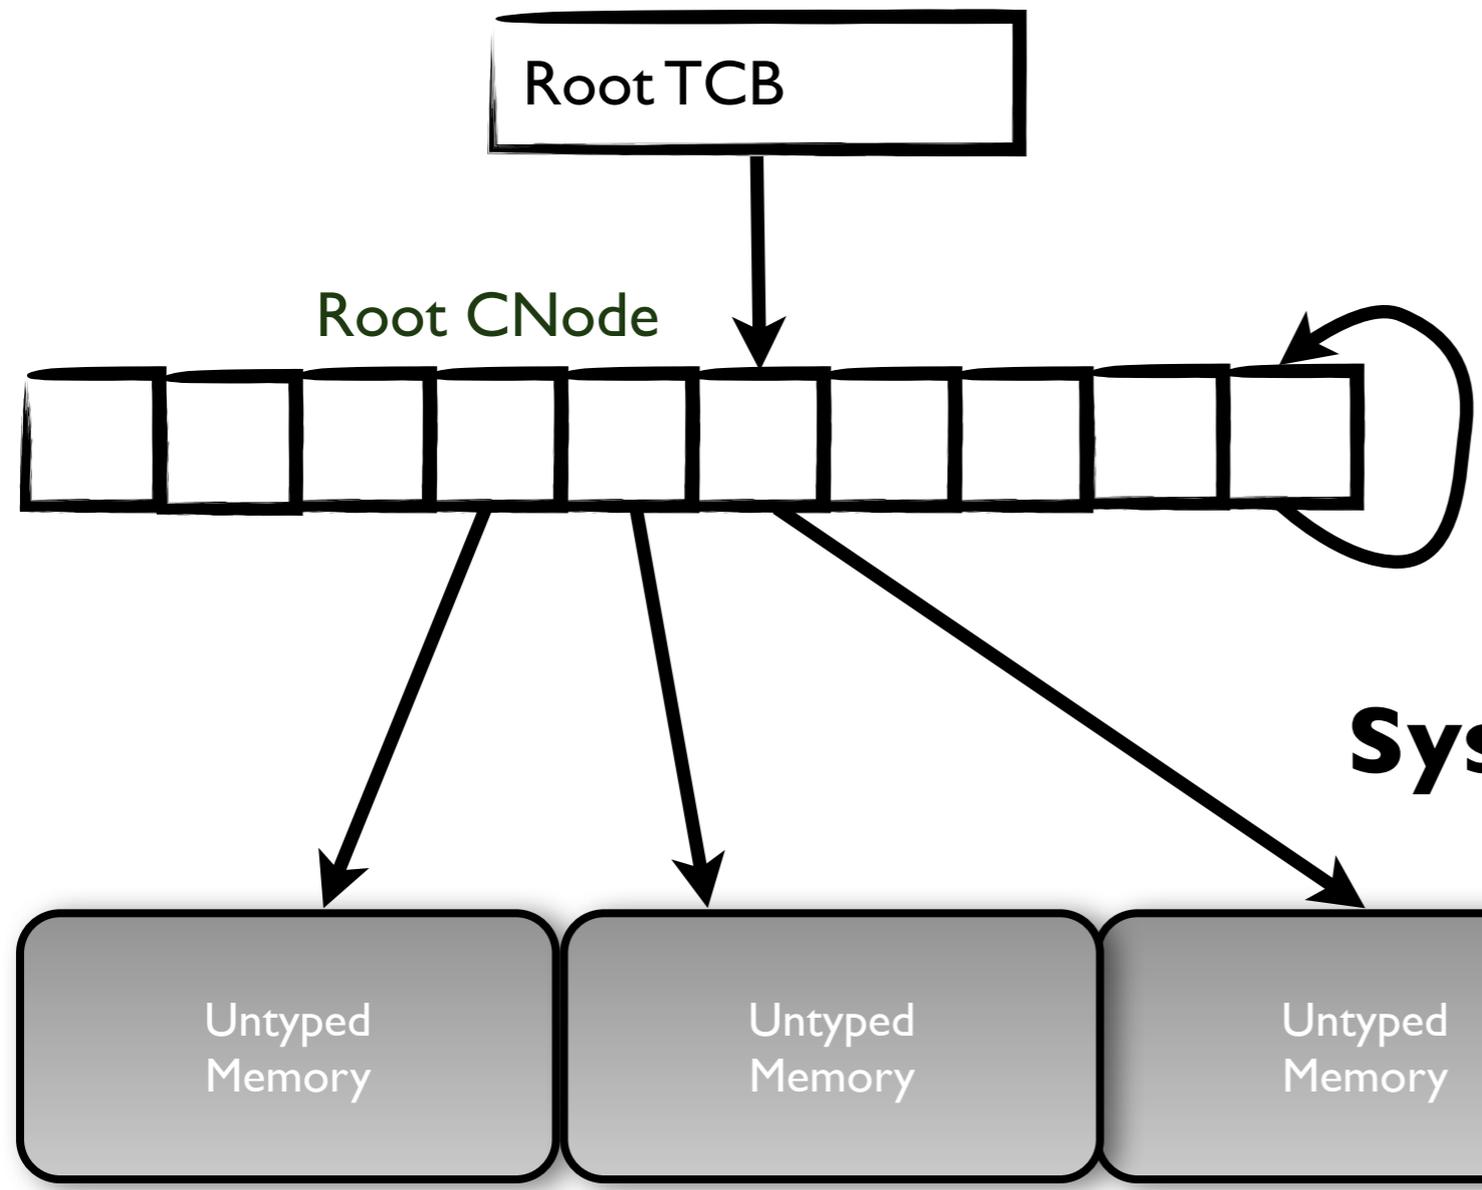
Initialised State

seL4



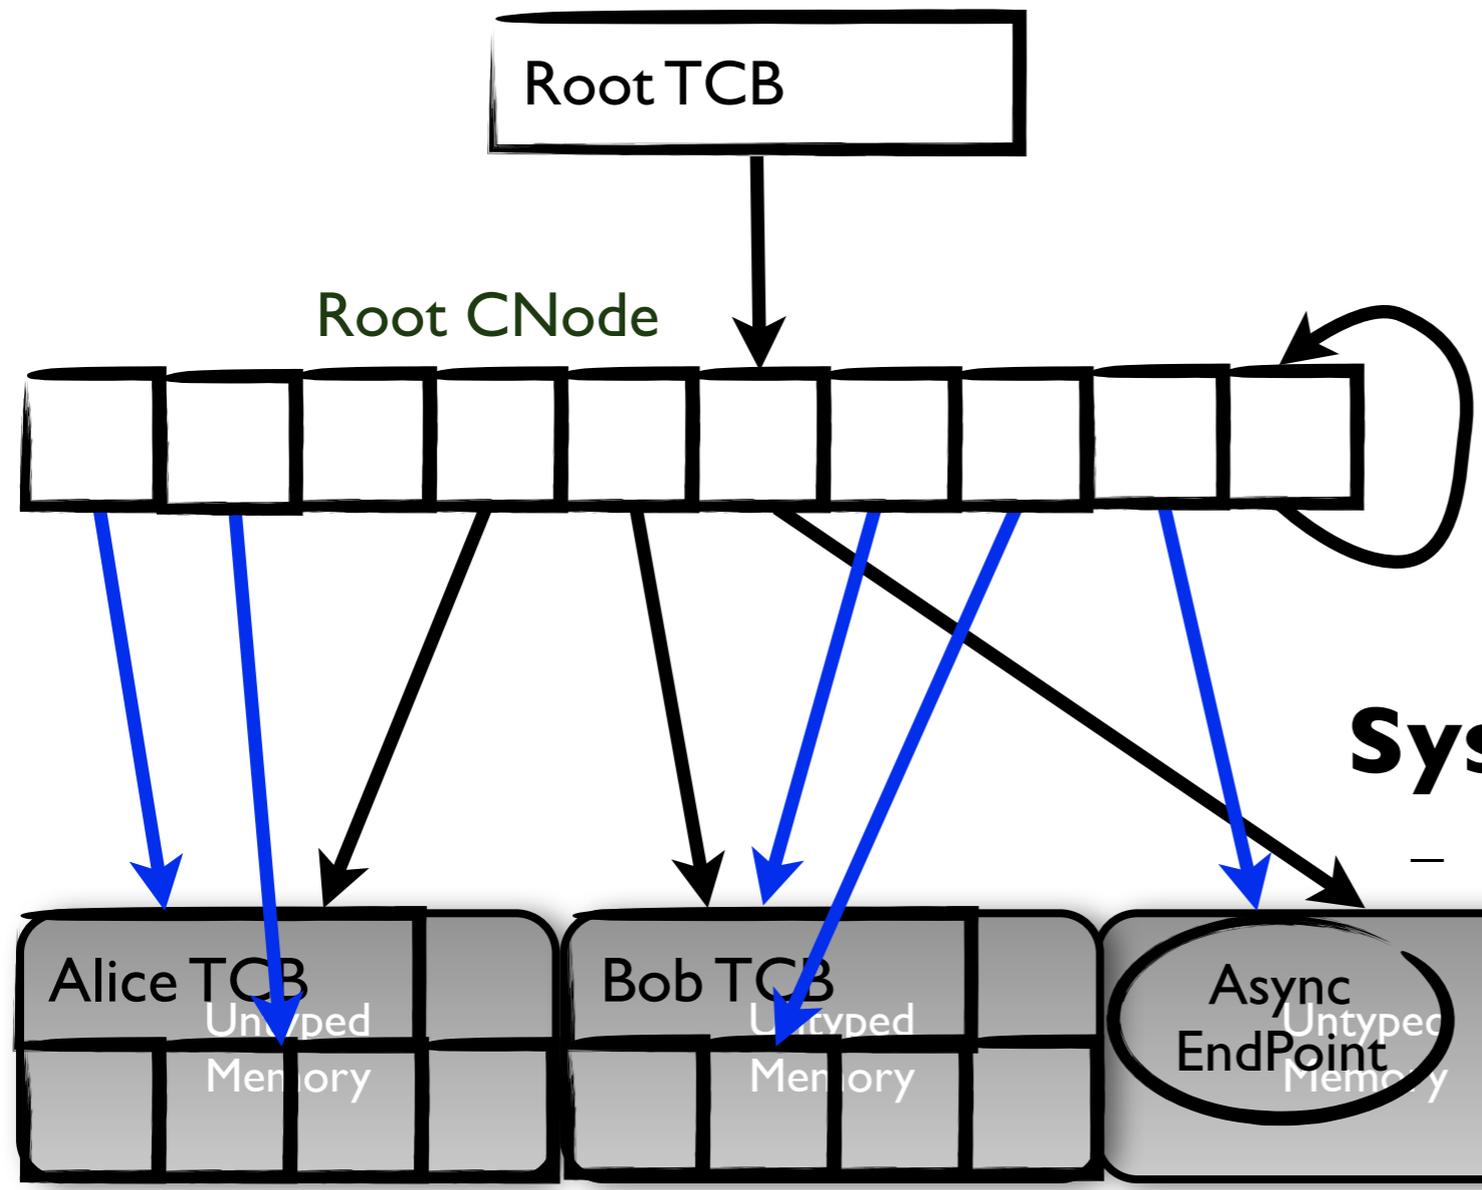
Hardware





System initialisation:



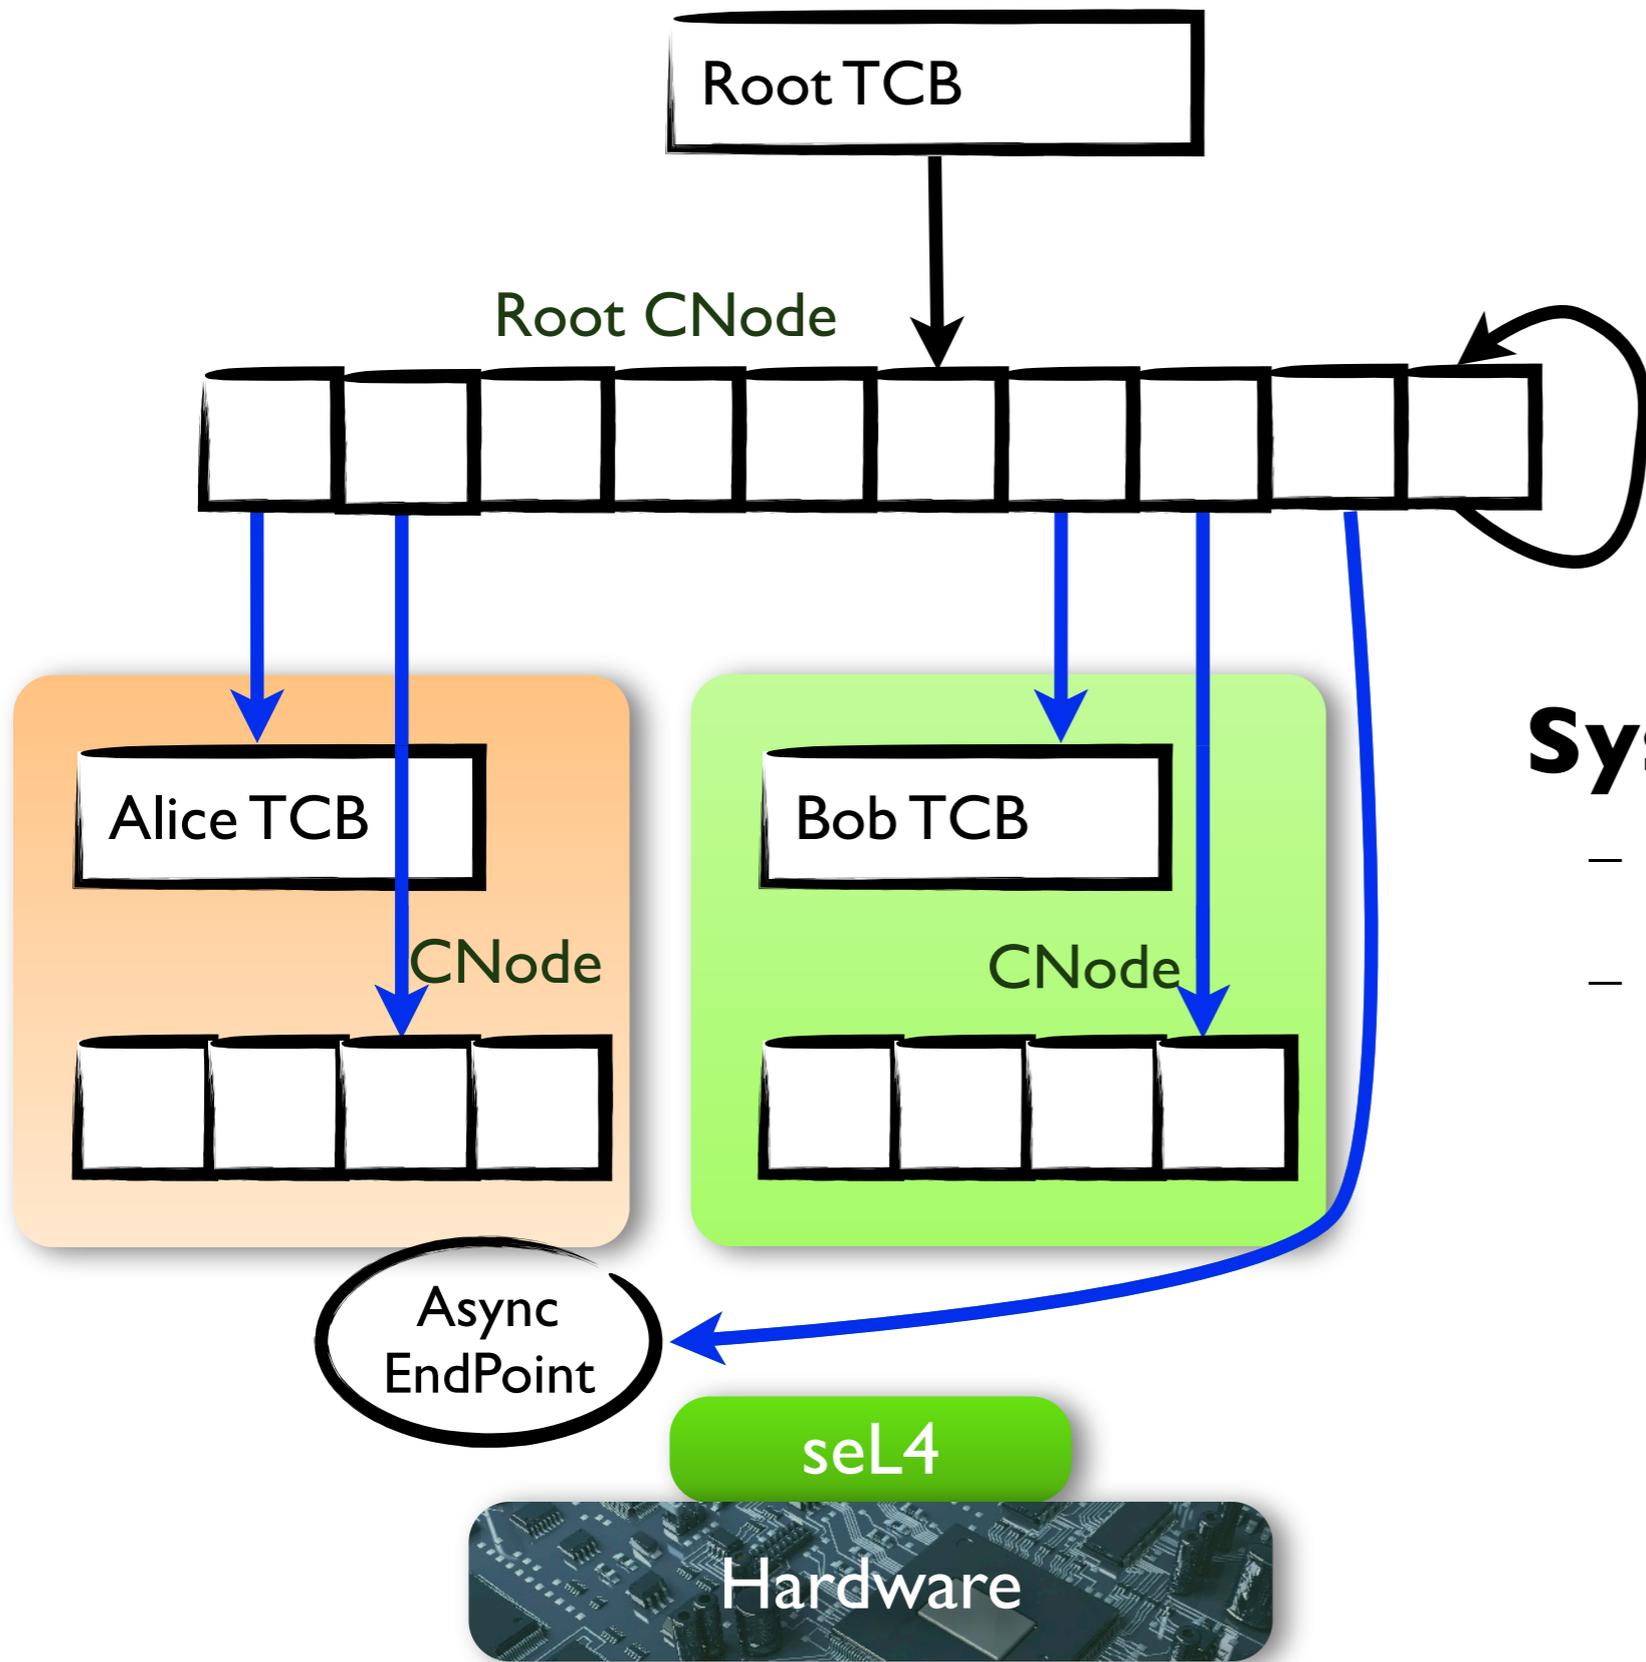


System initialisation:

- Create objects

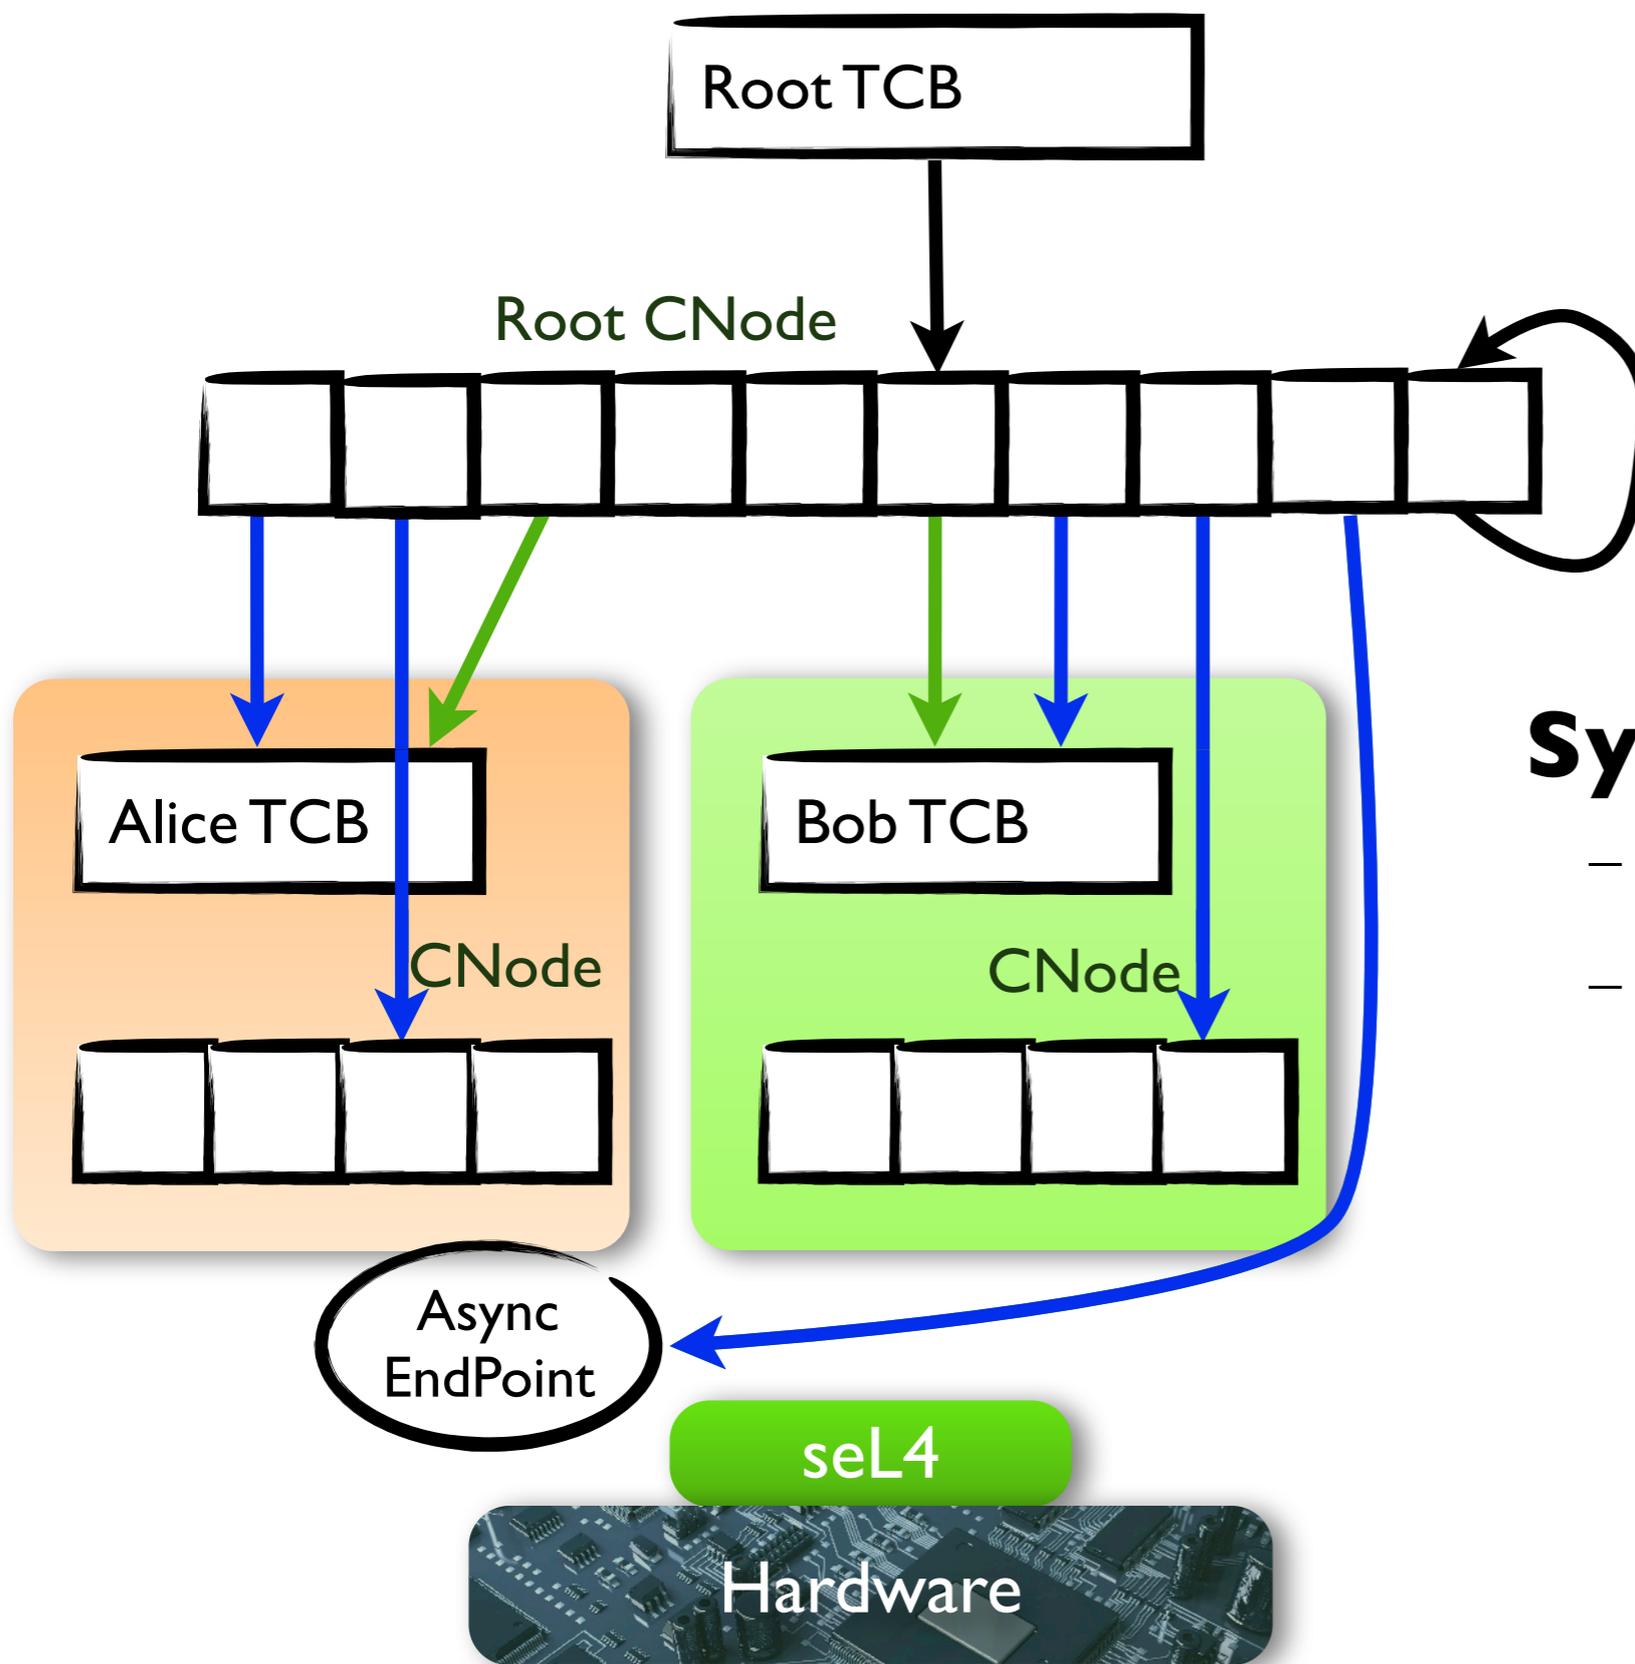
seL4





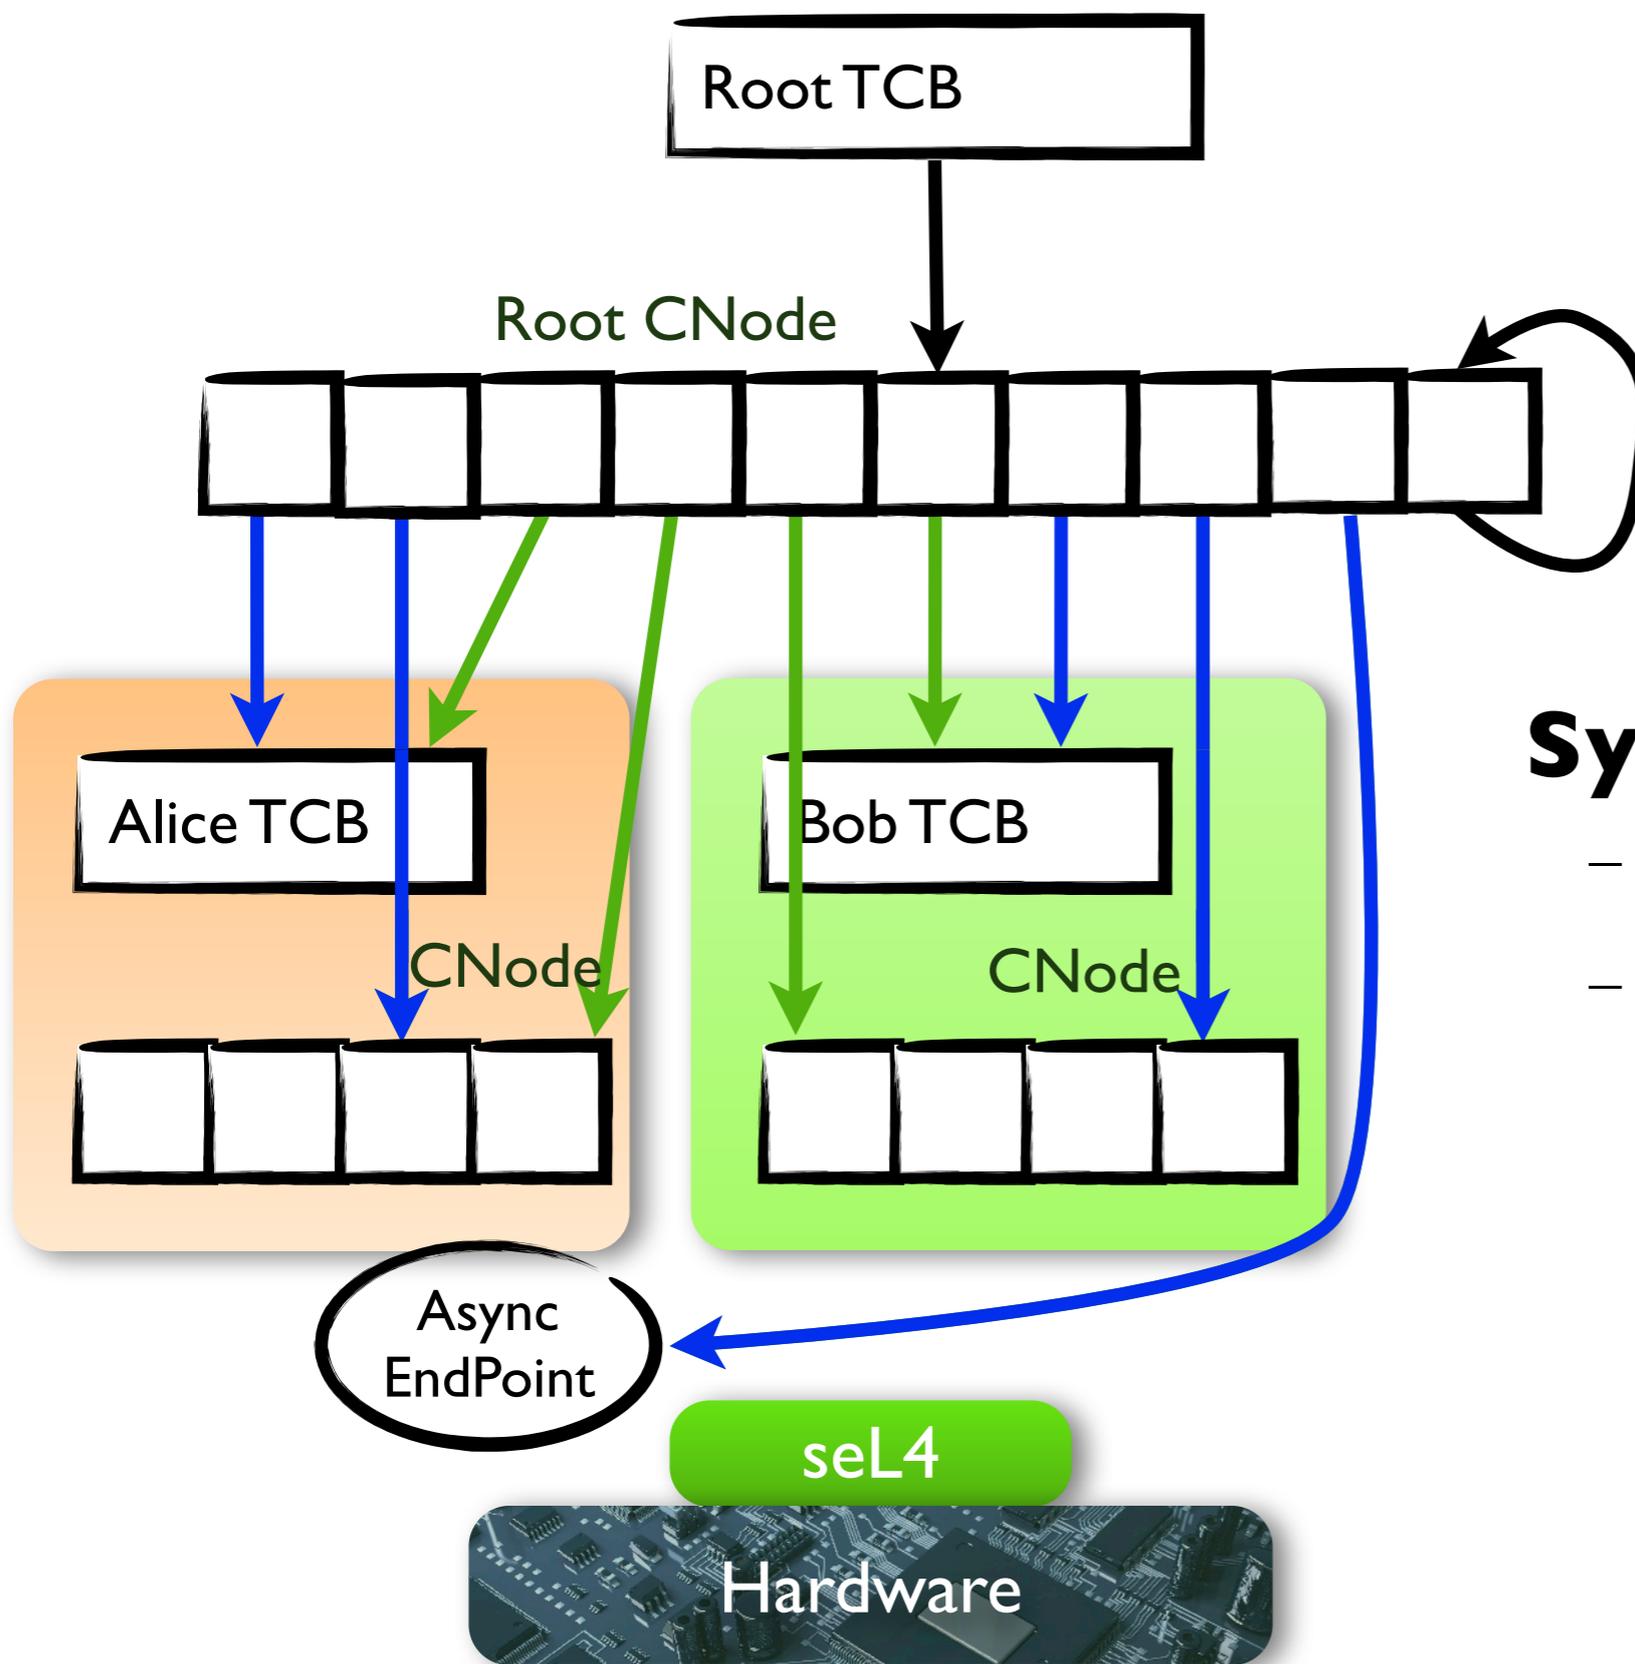
System initialisation:

- Create objects
- Duplicate capabilities



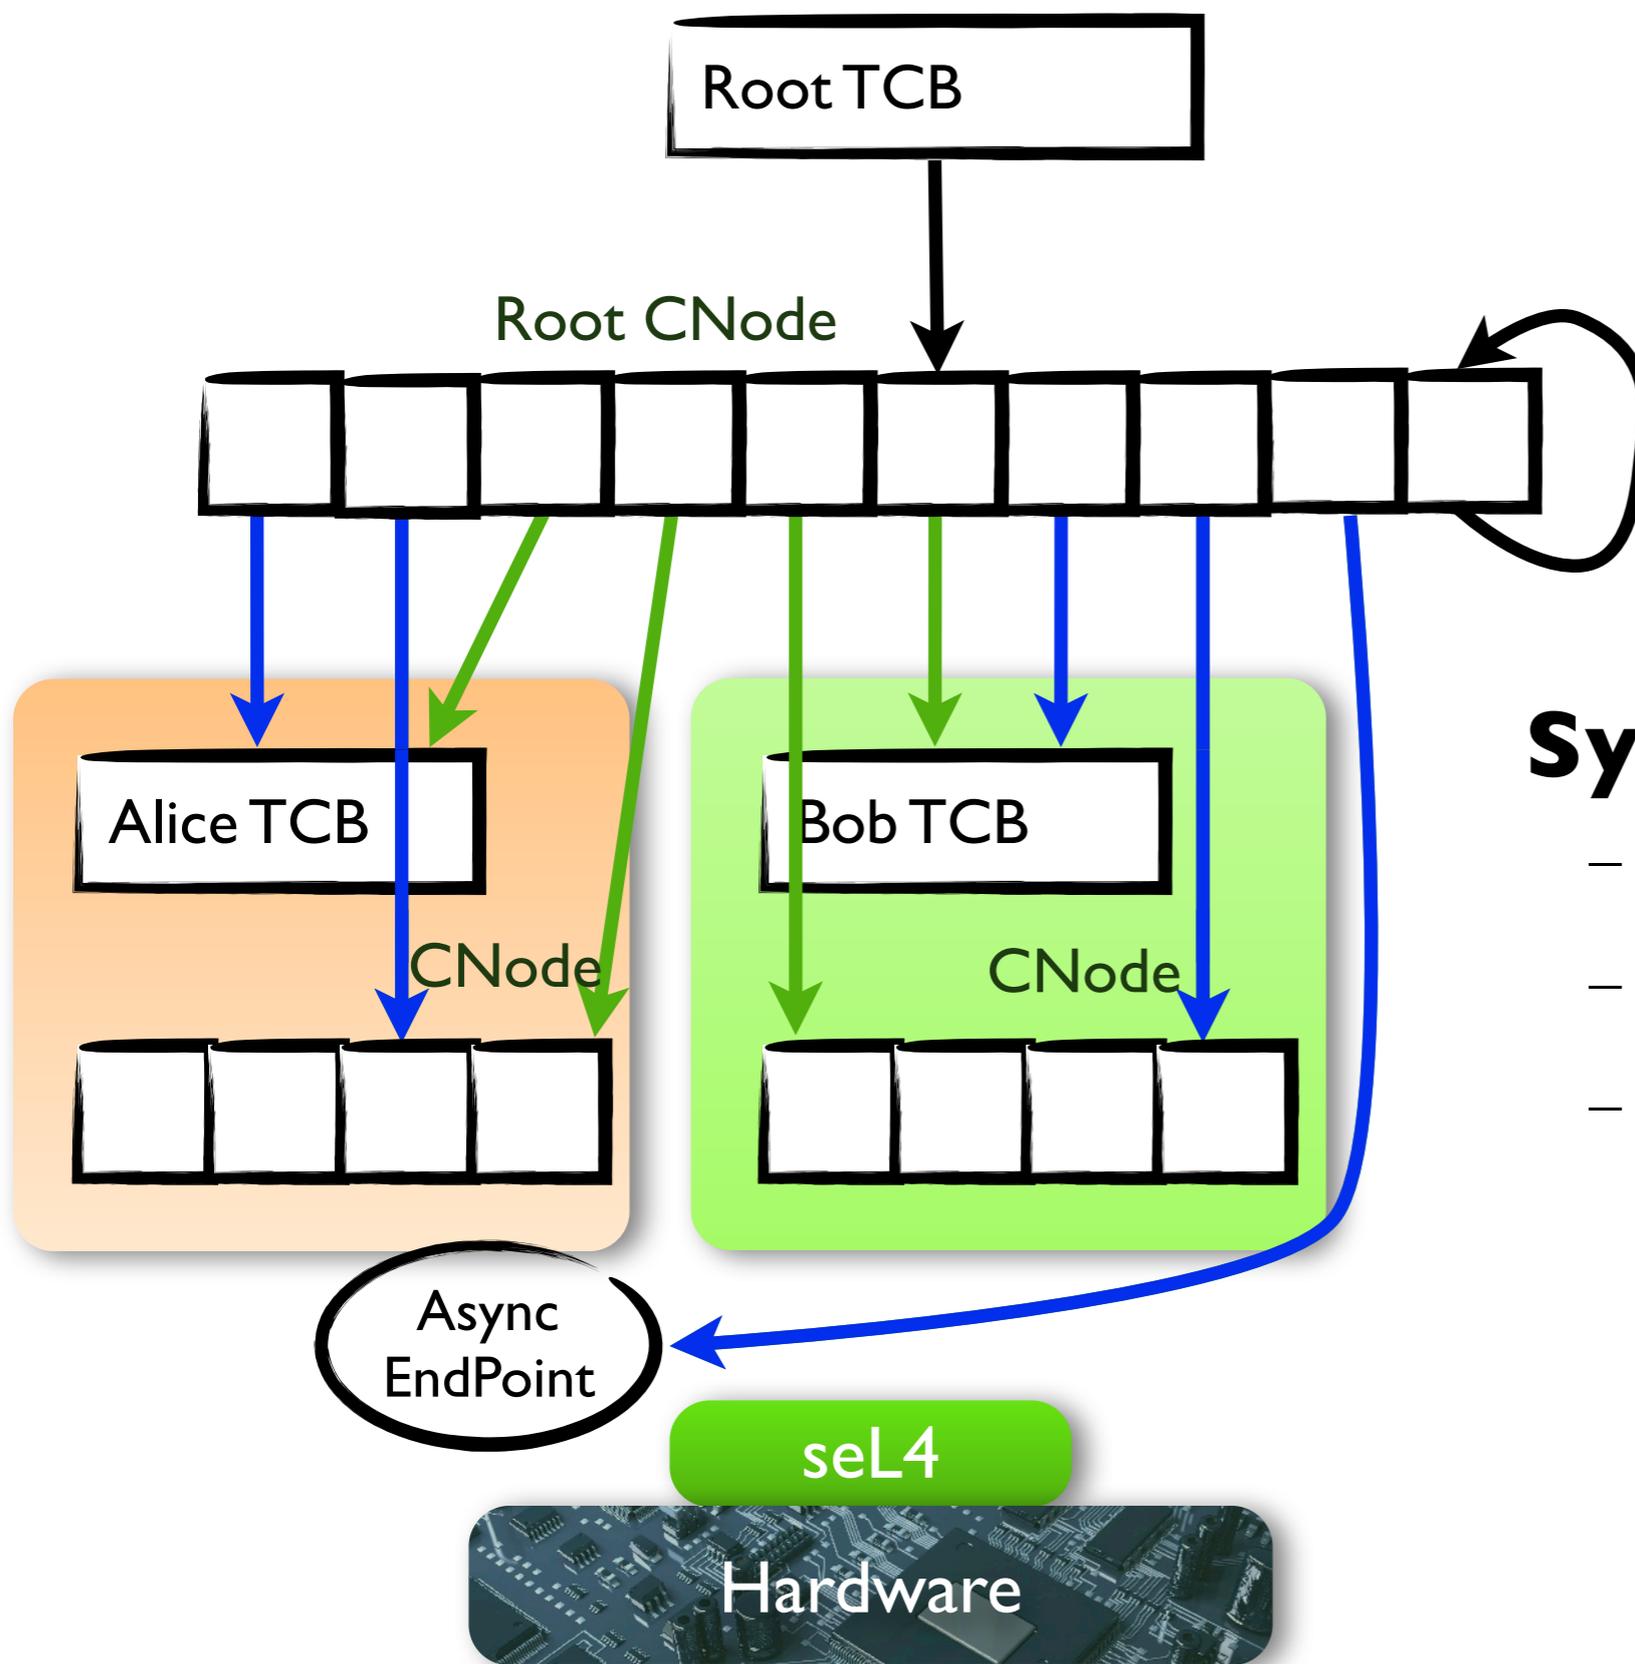
System initialisation:

- Create objects
- Duplicate capabilities



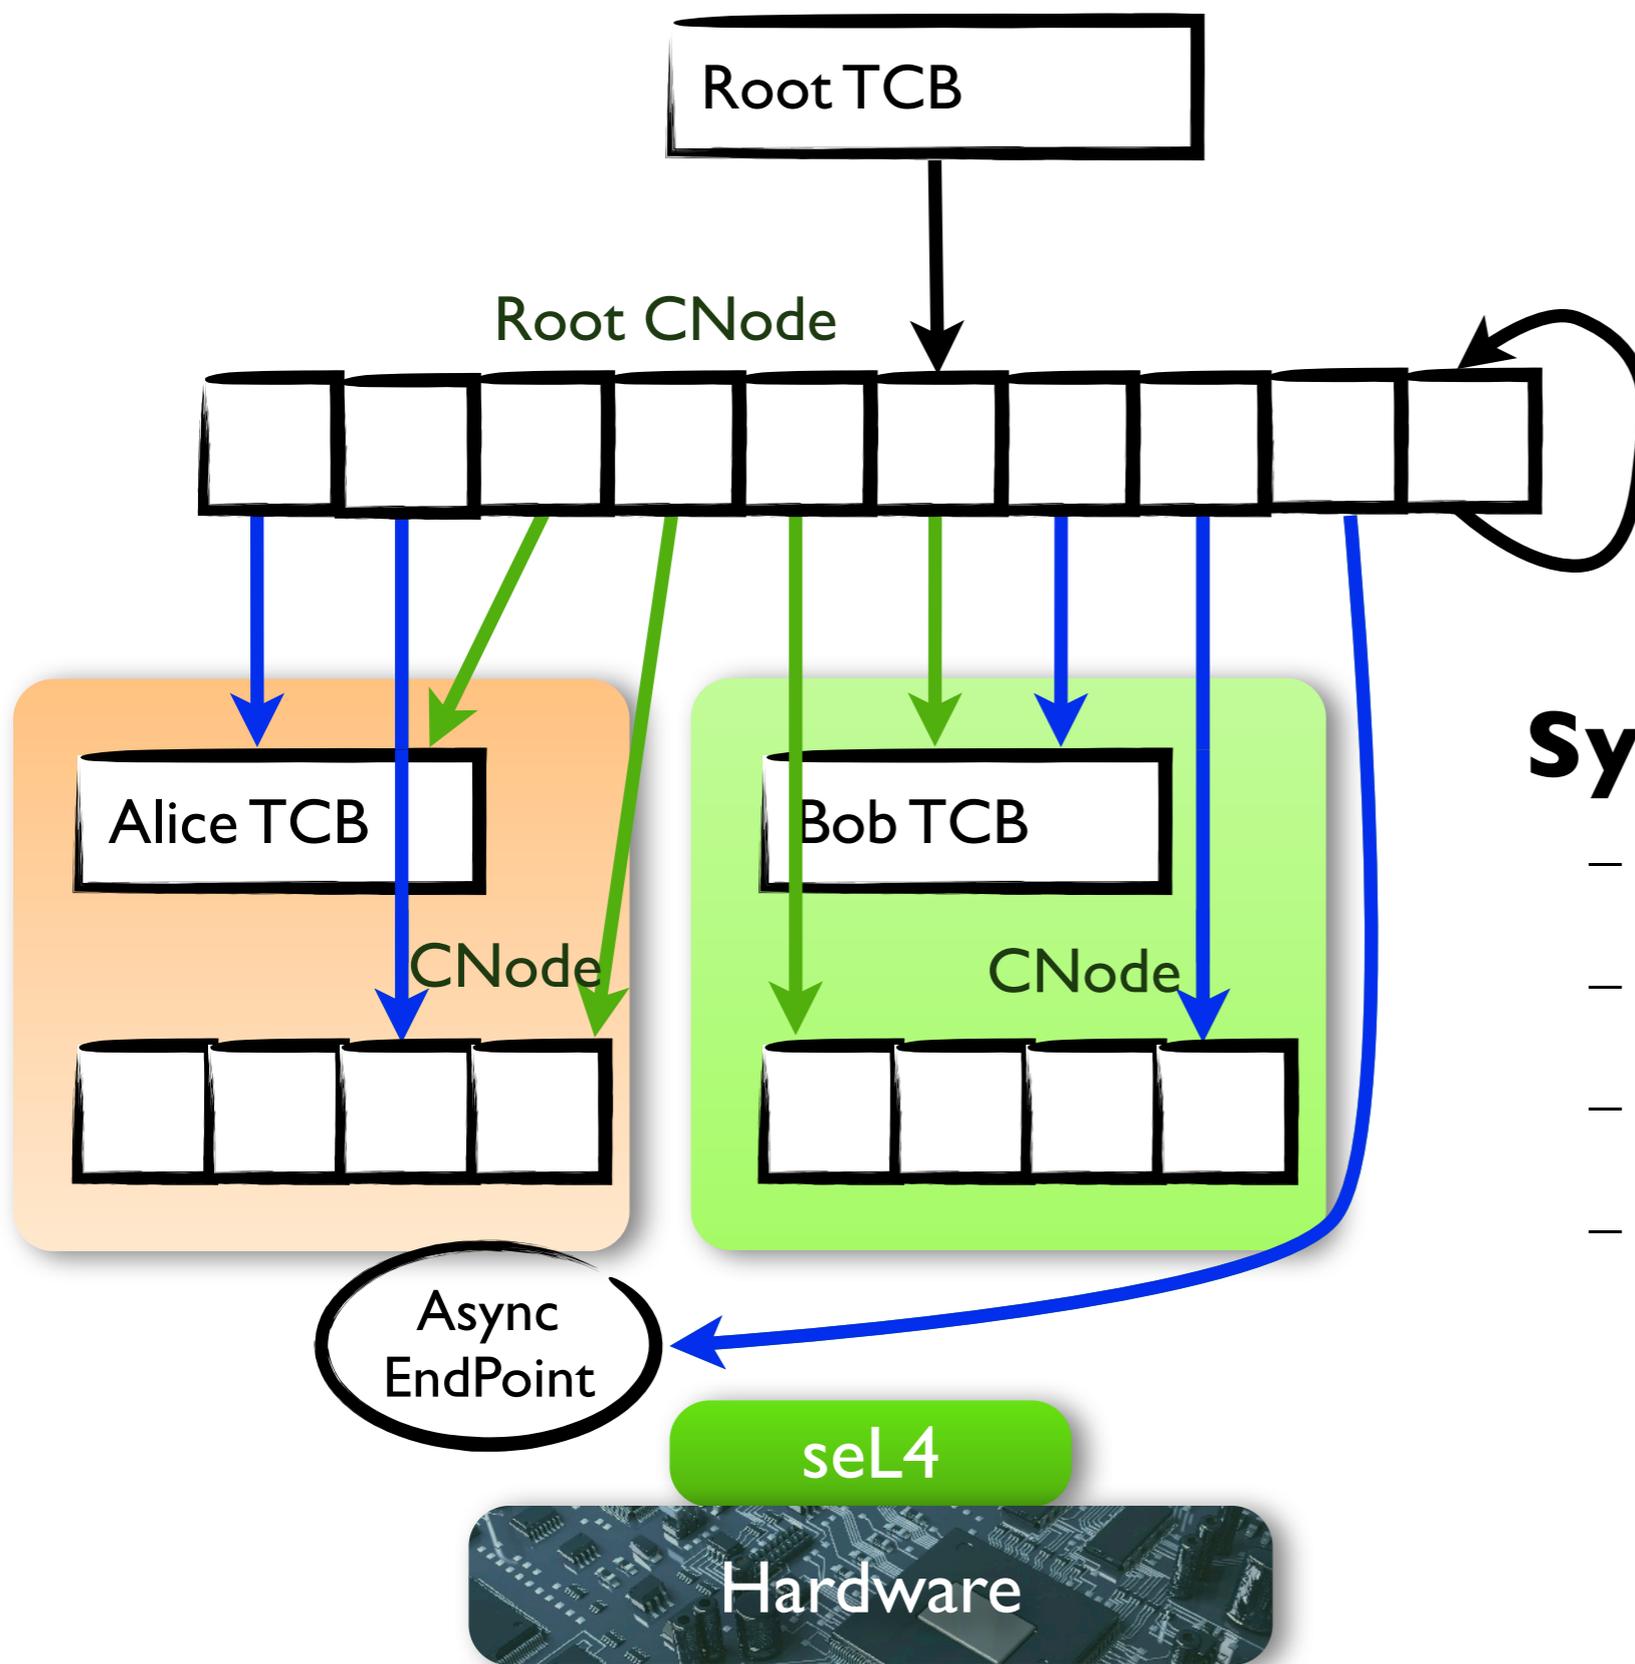
System initialisation:

- Create objects
- Duplicate capabilities



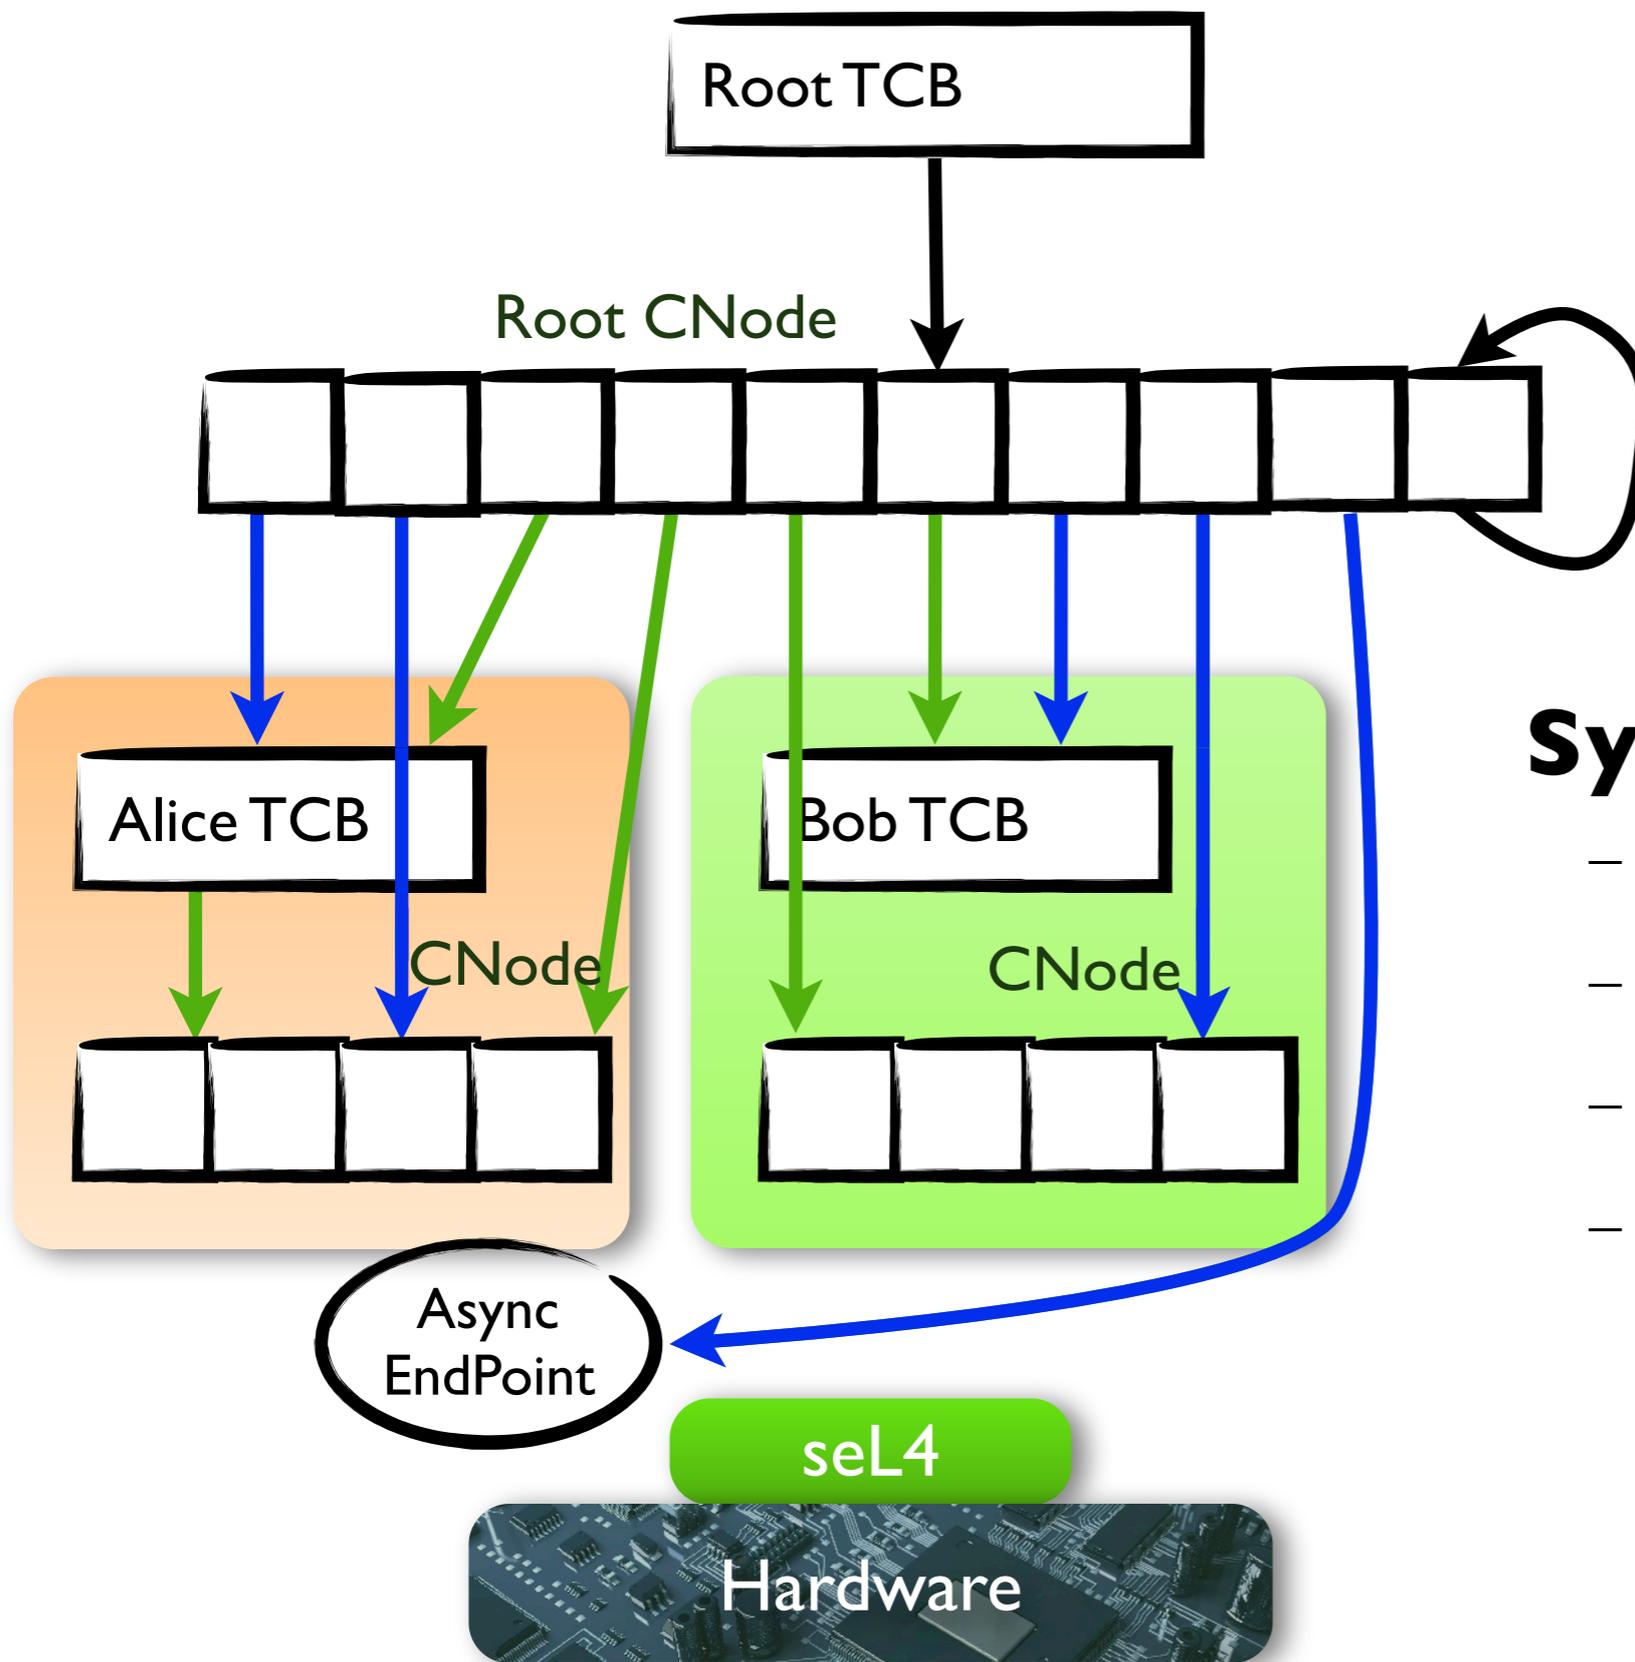
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space



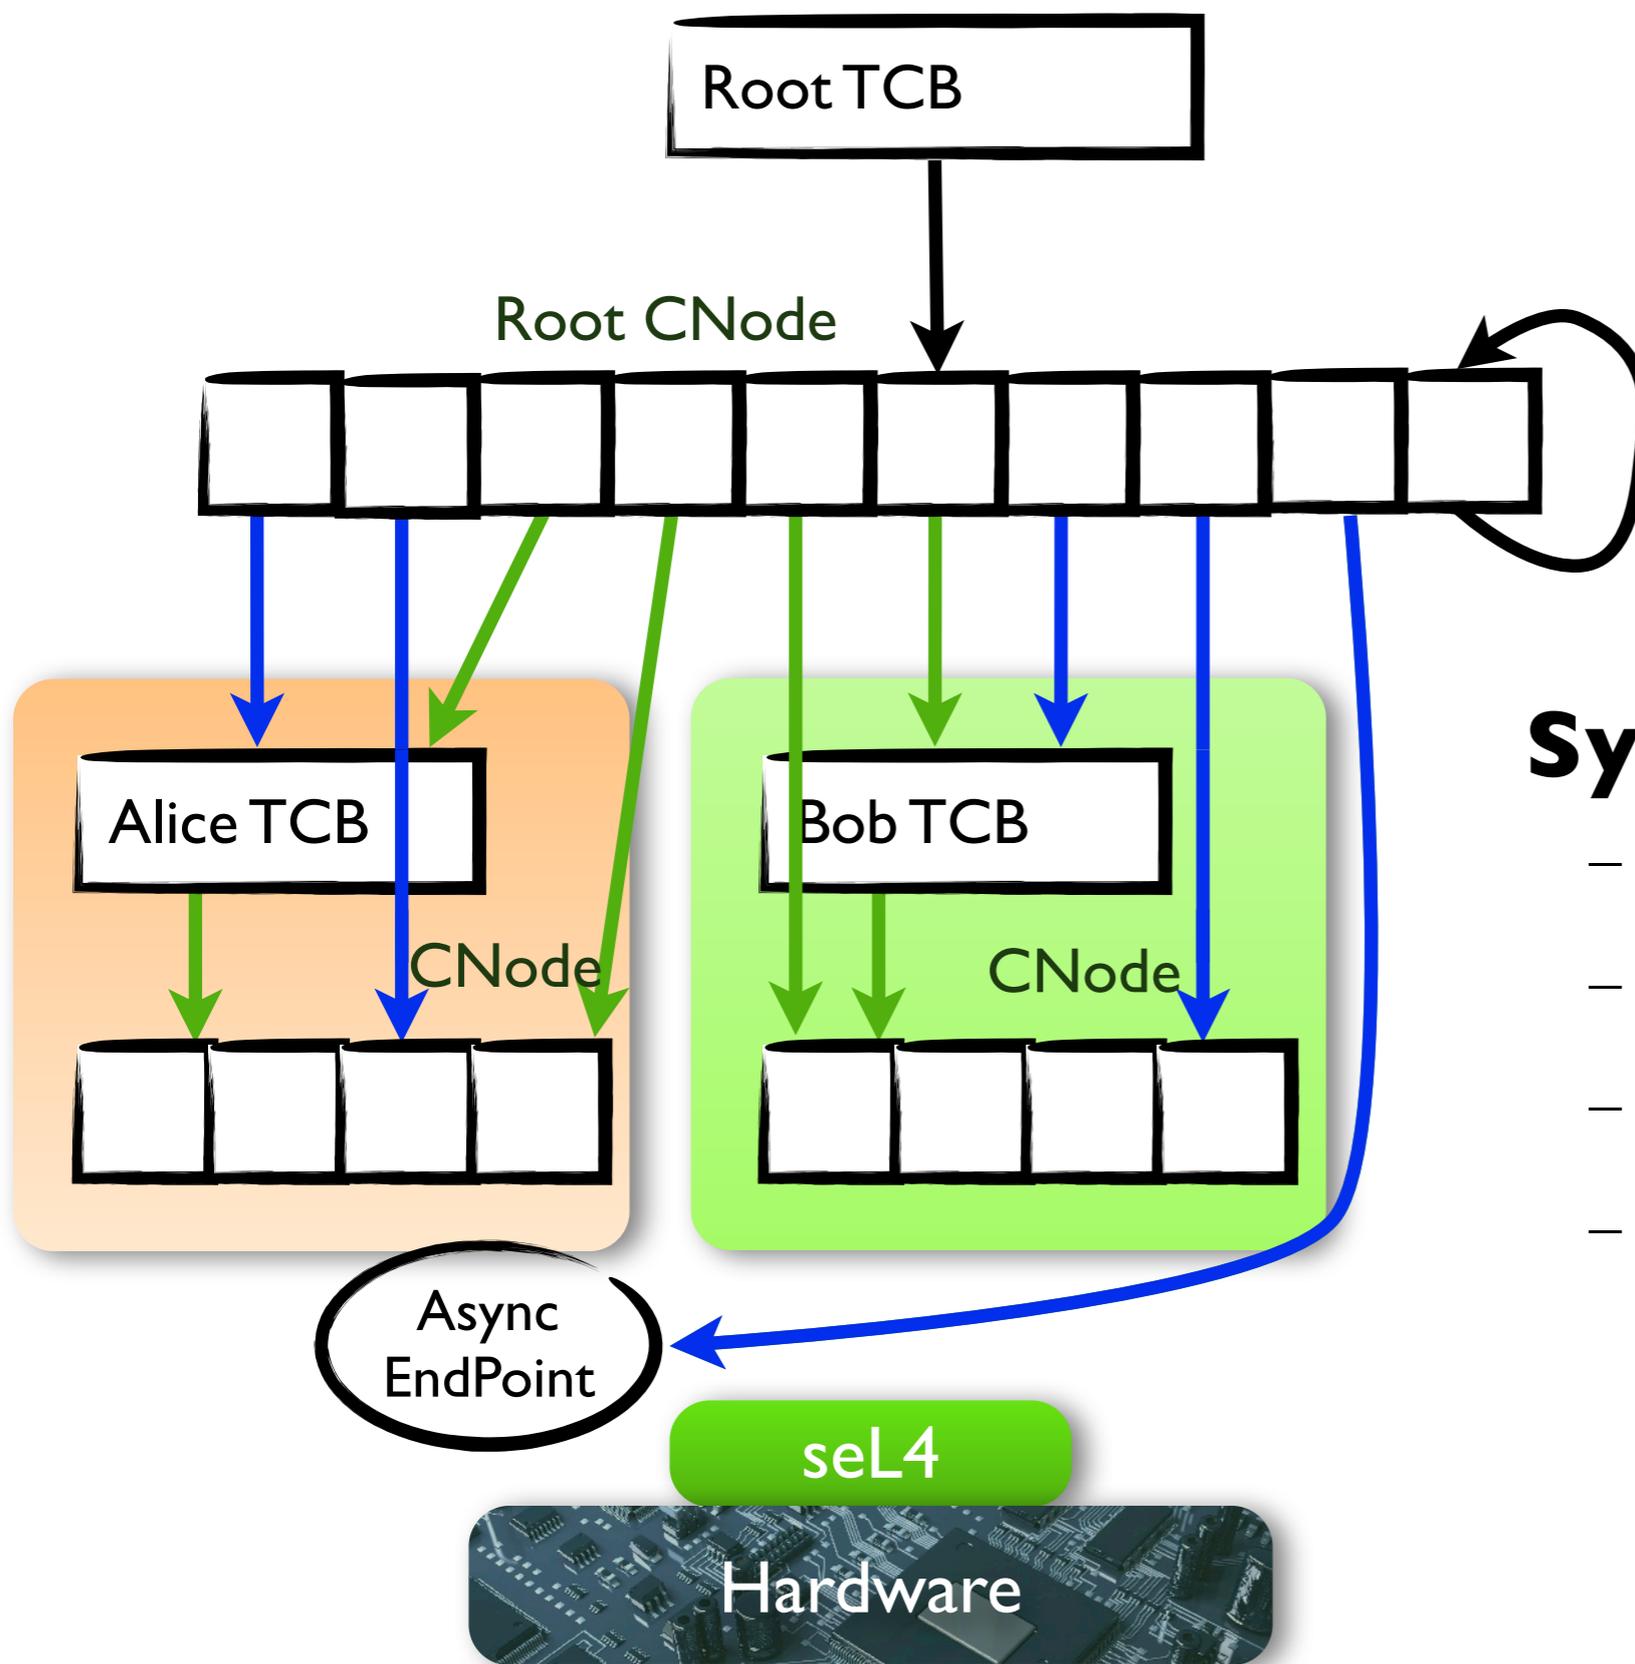
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks



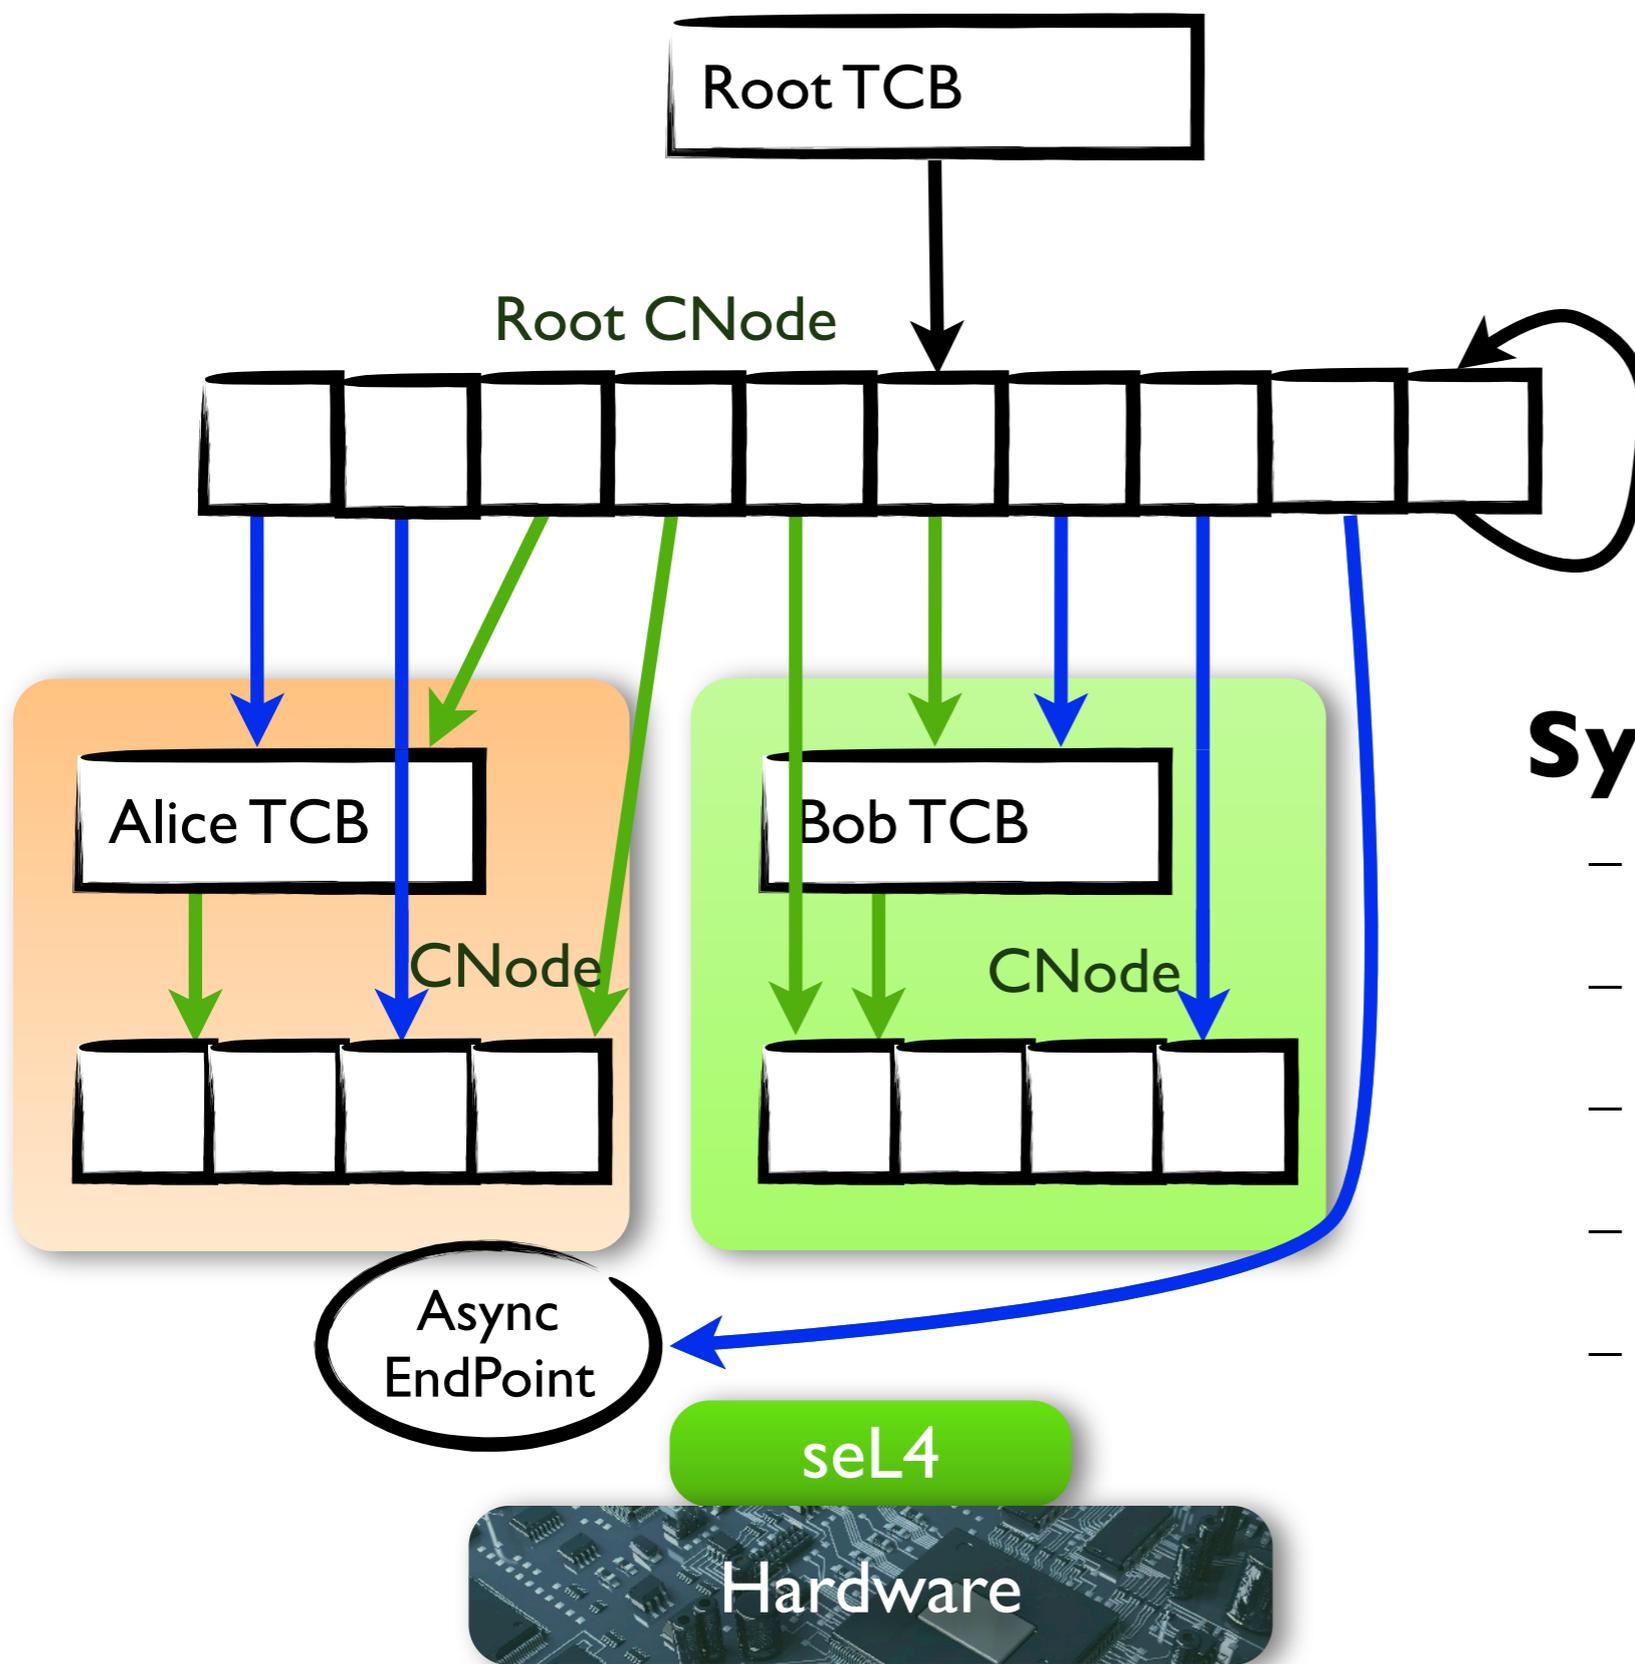
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks



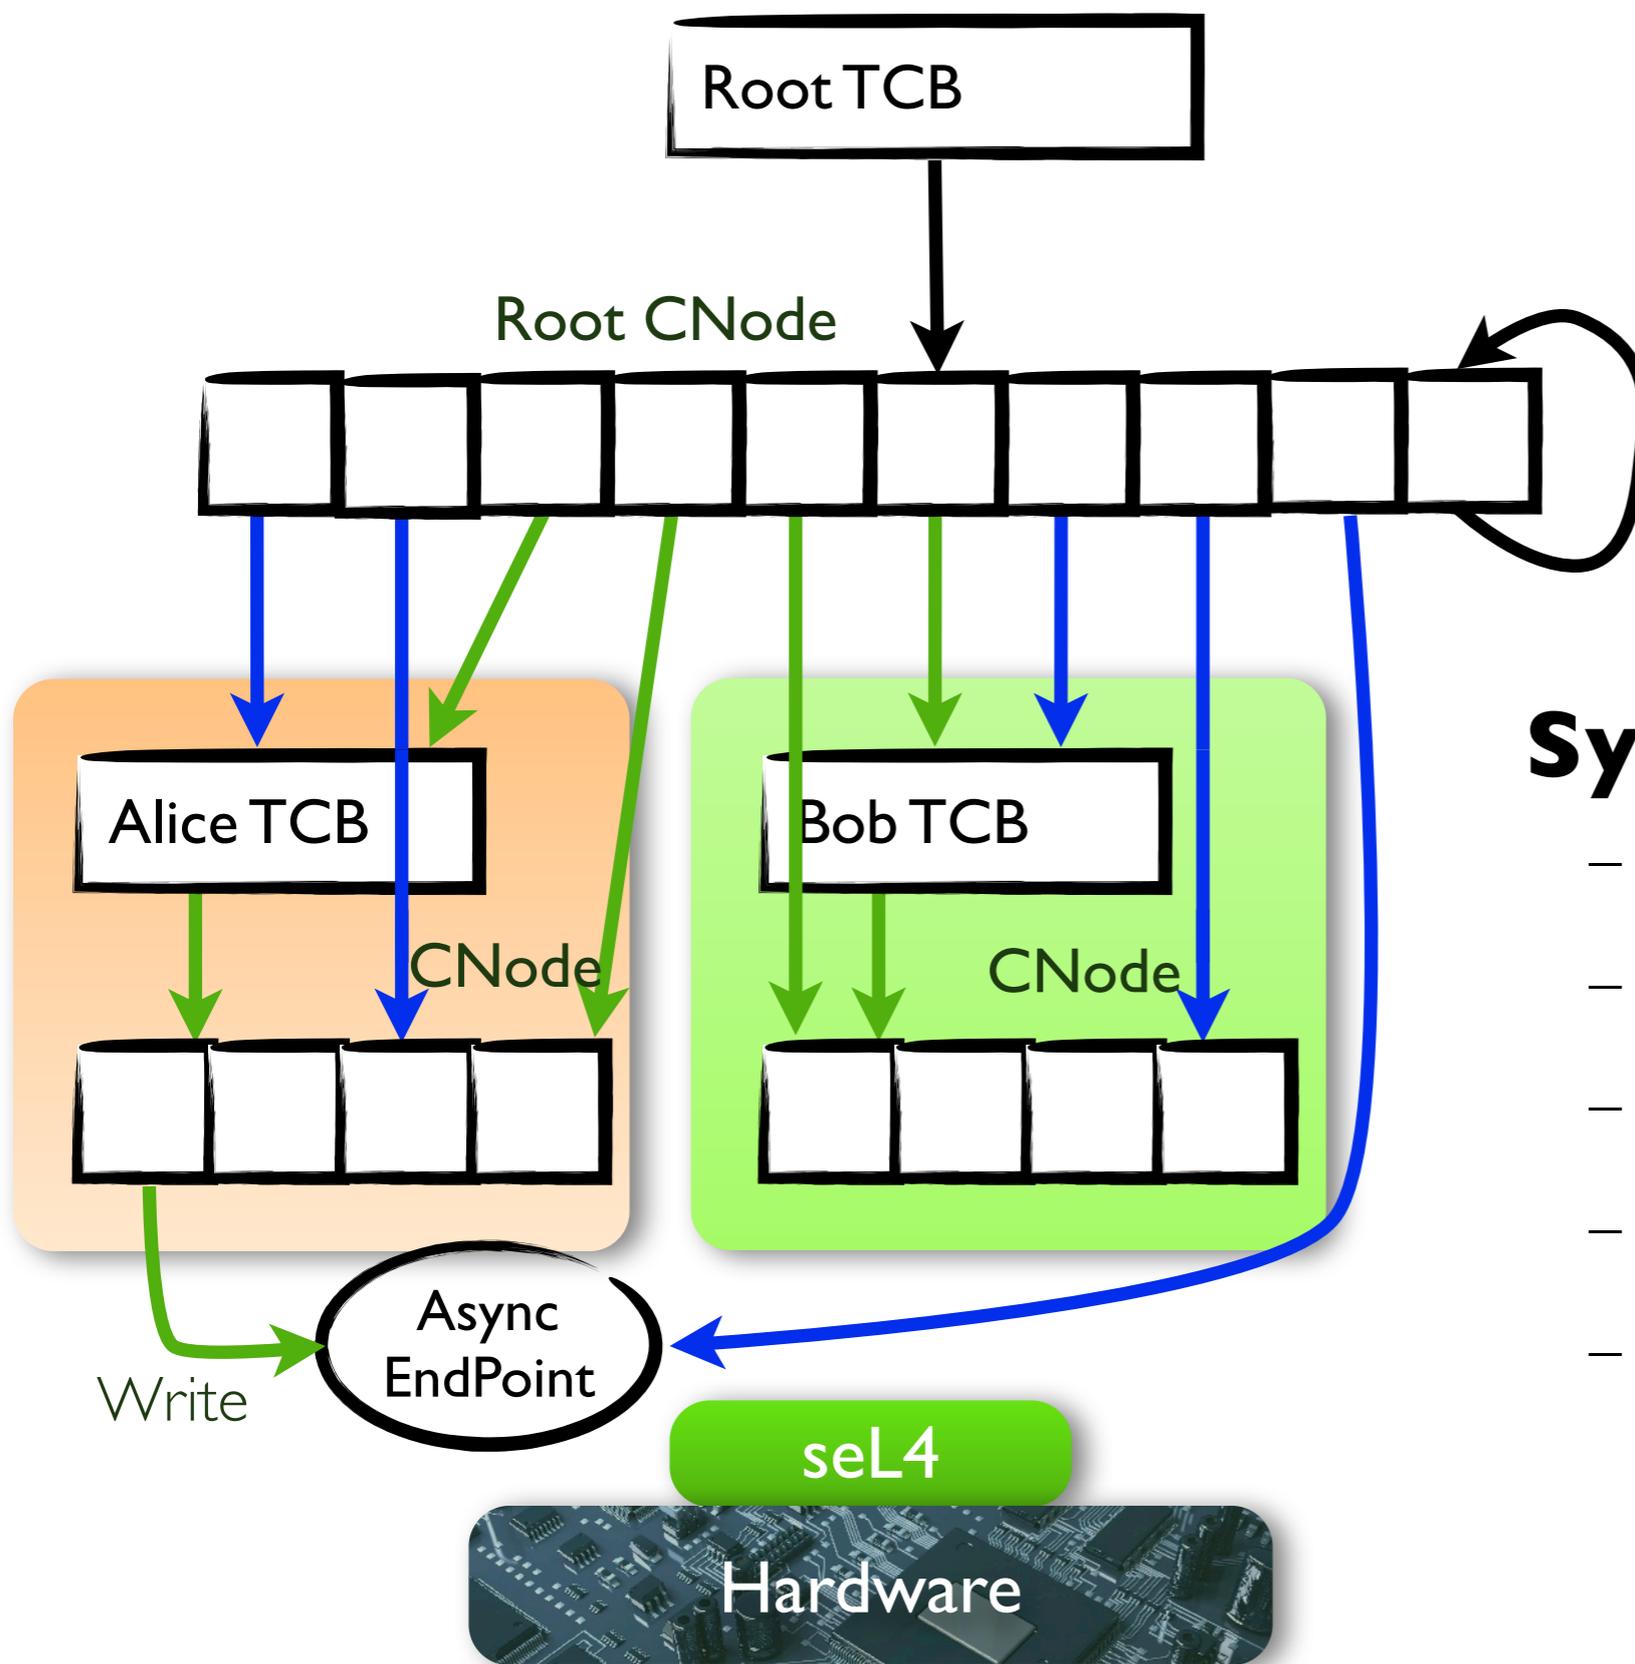
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks



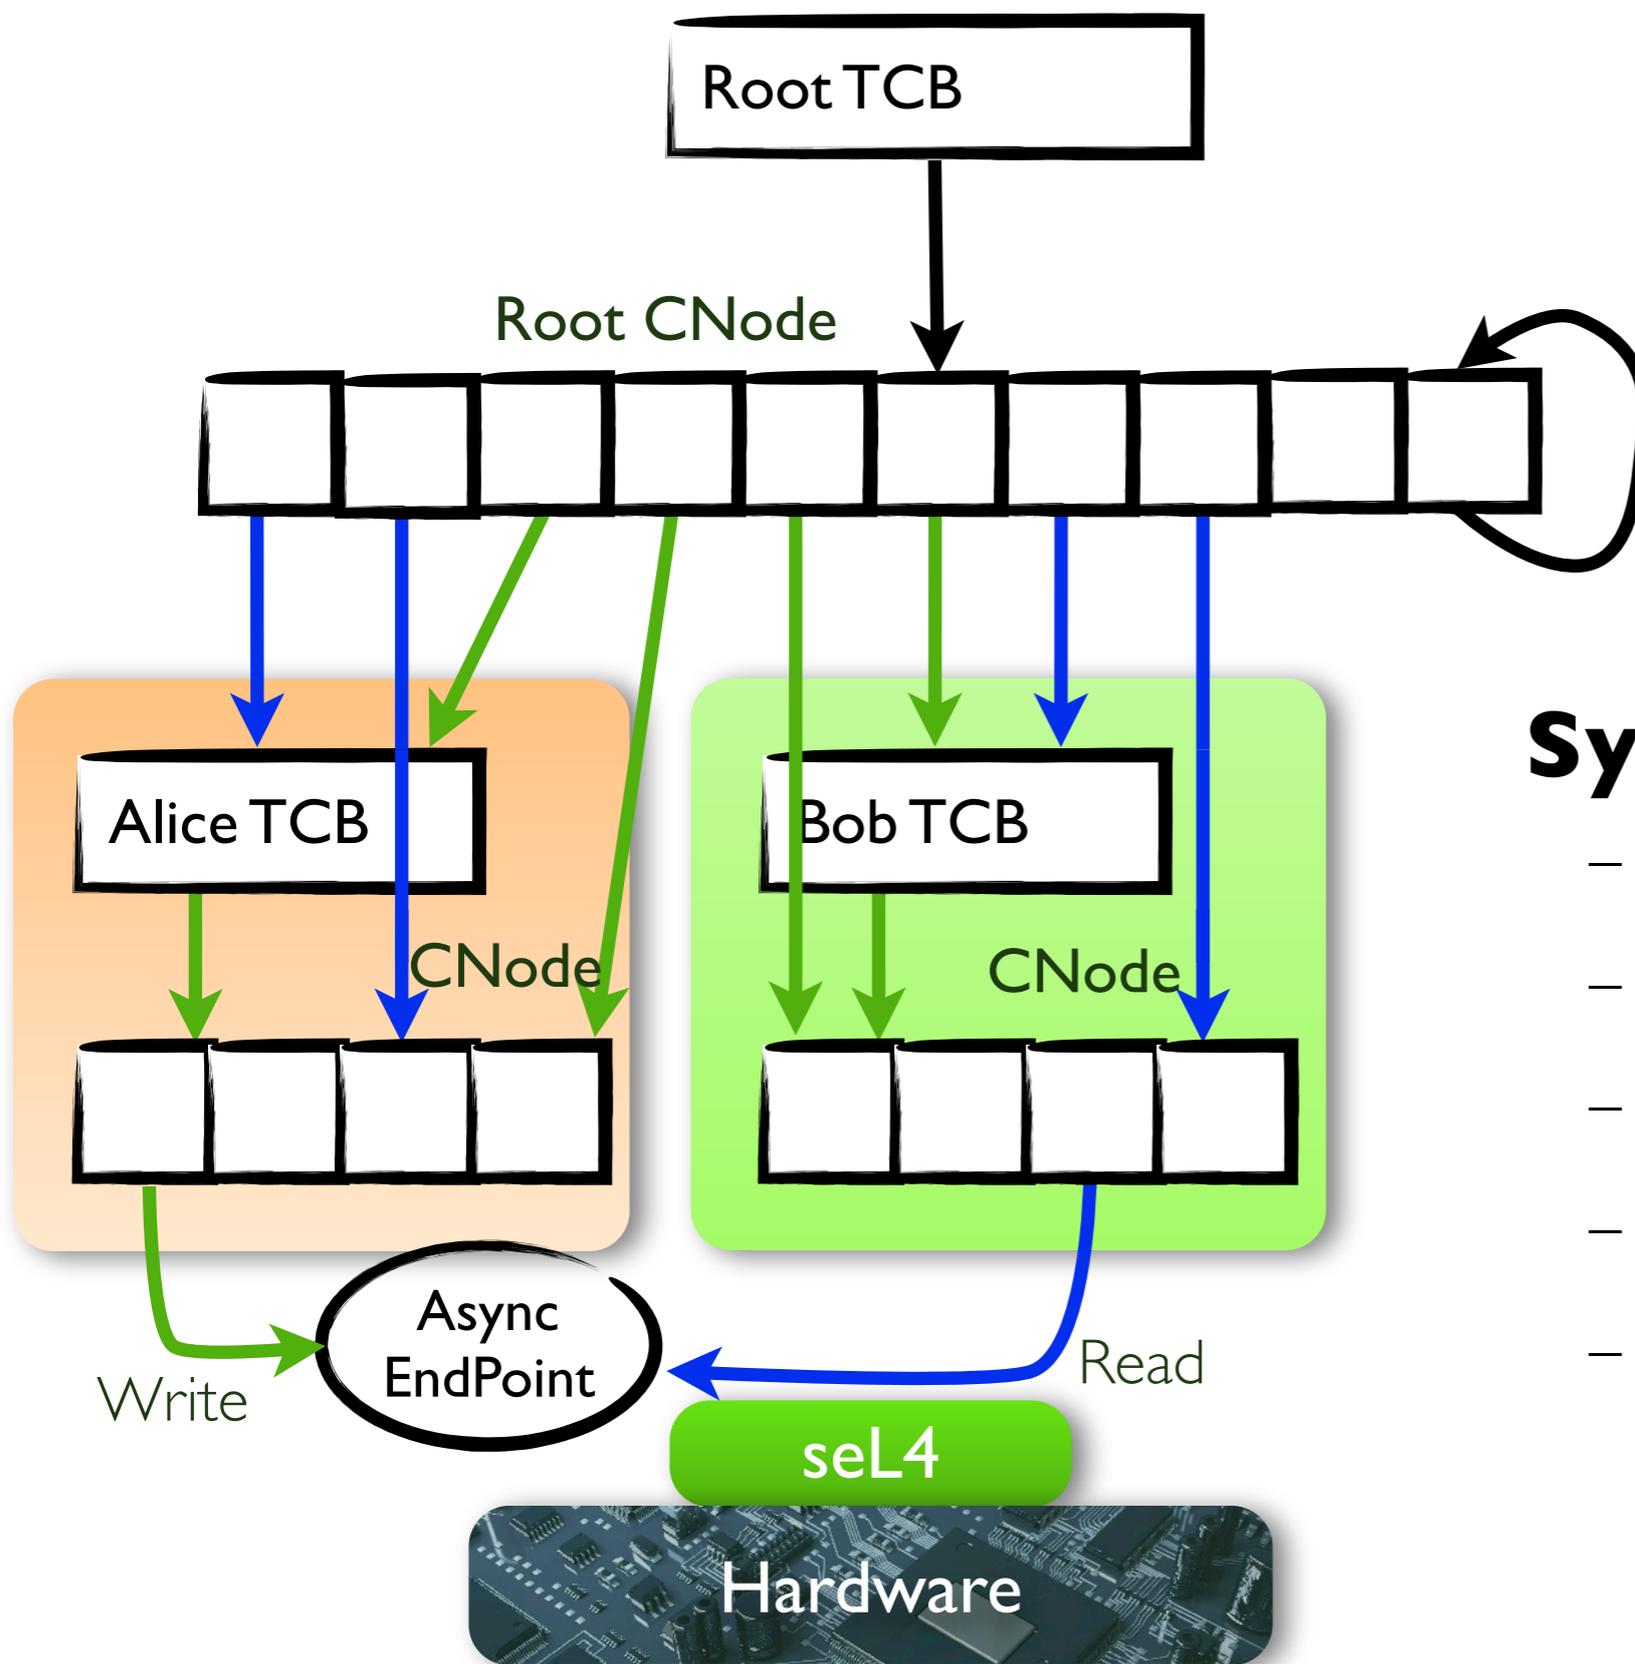
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



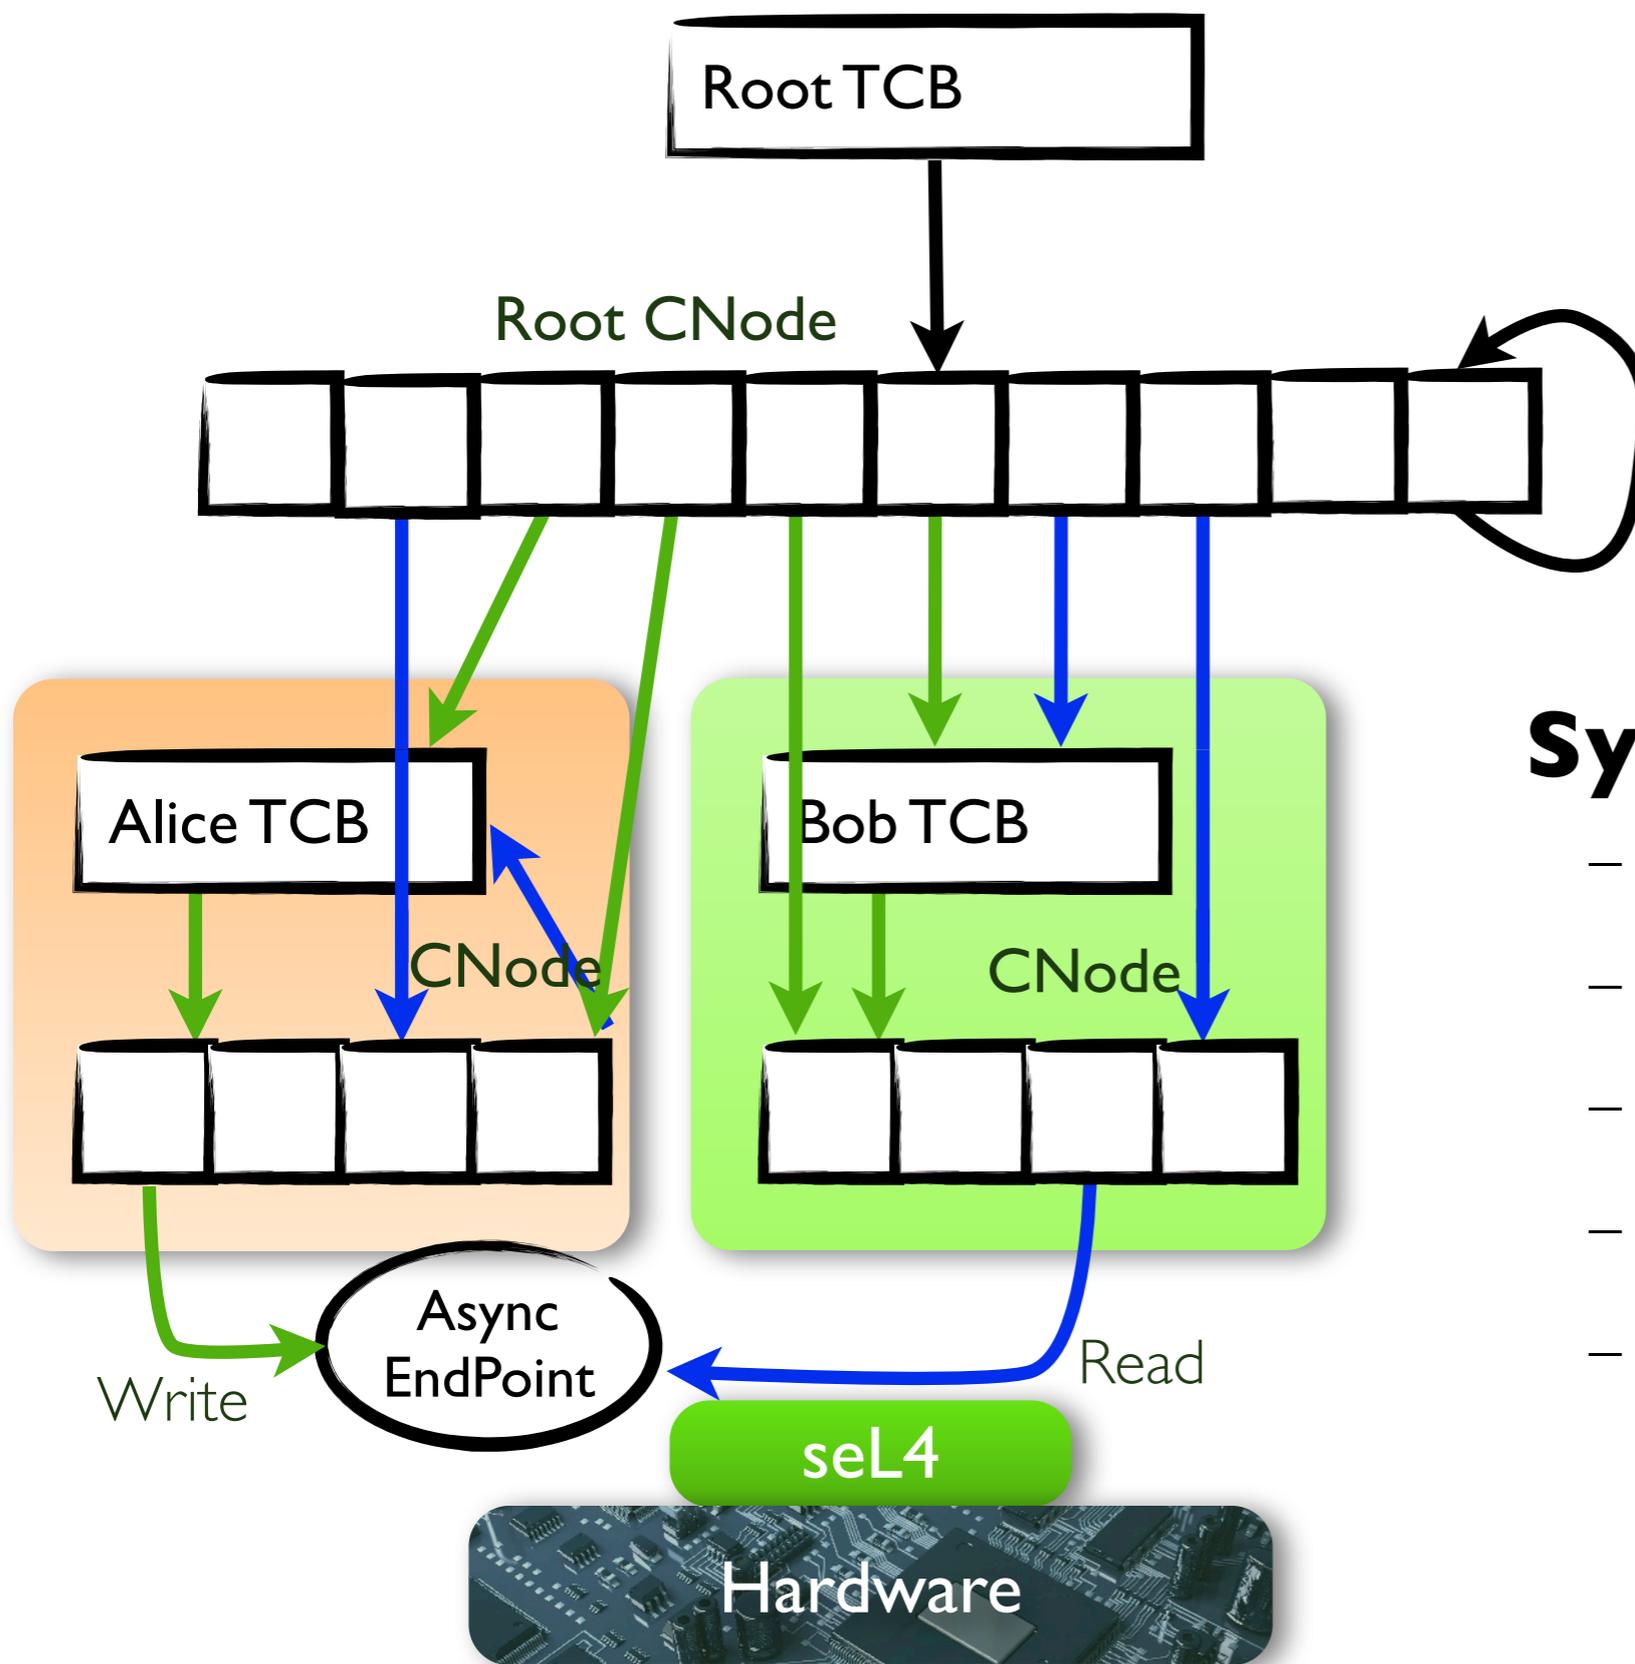
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



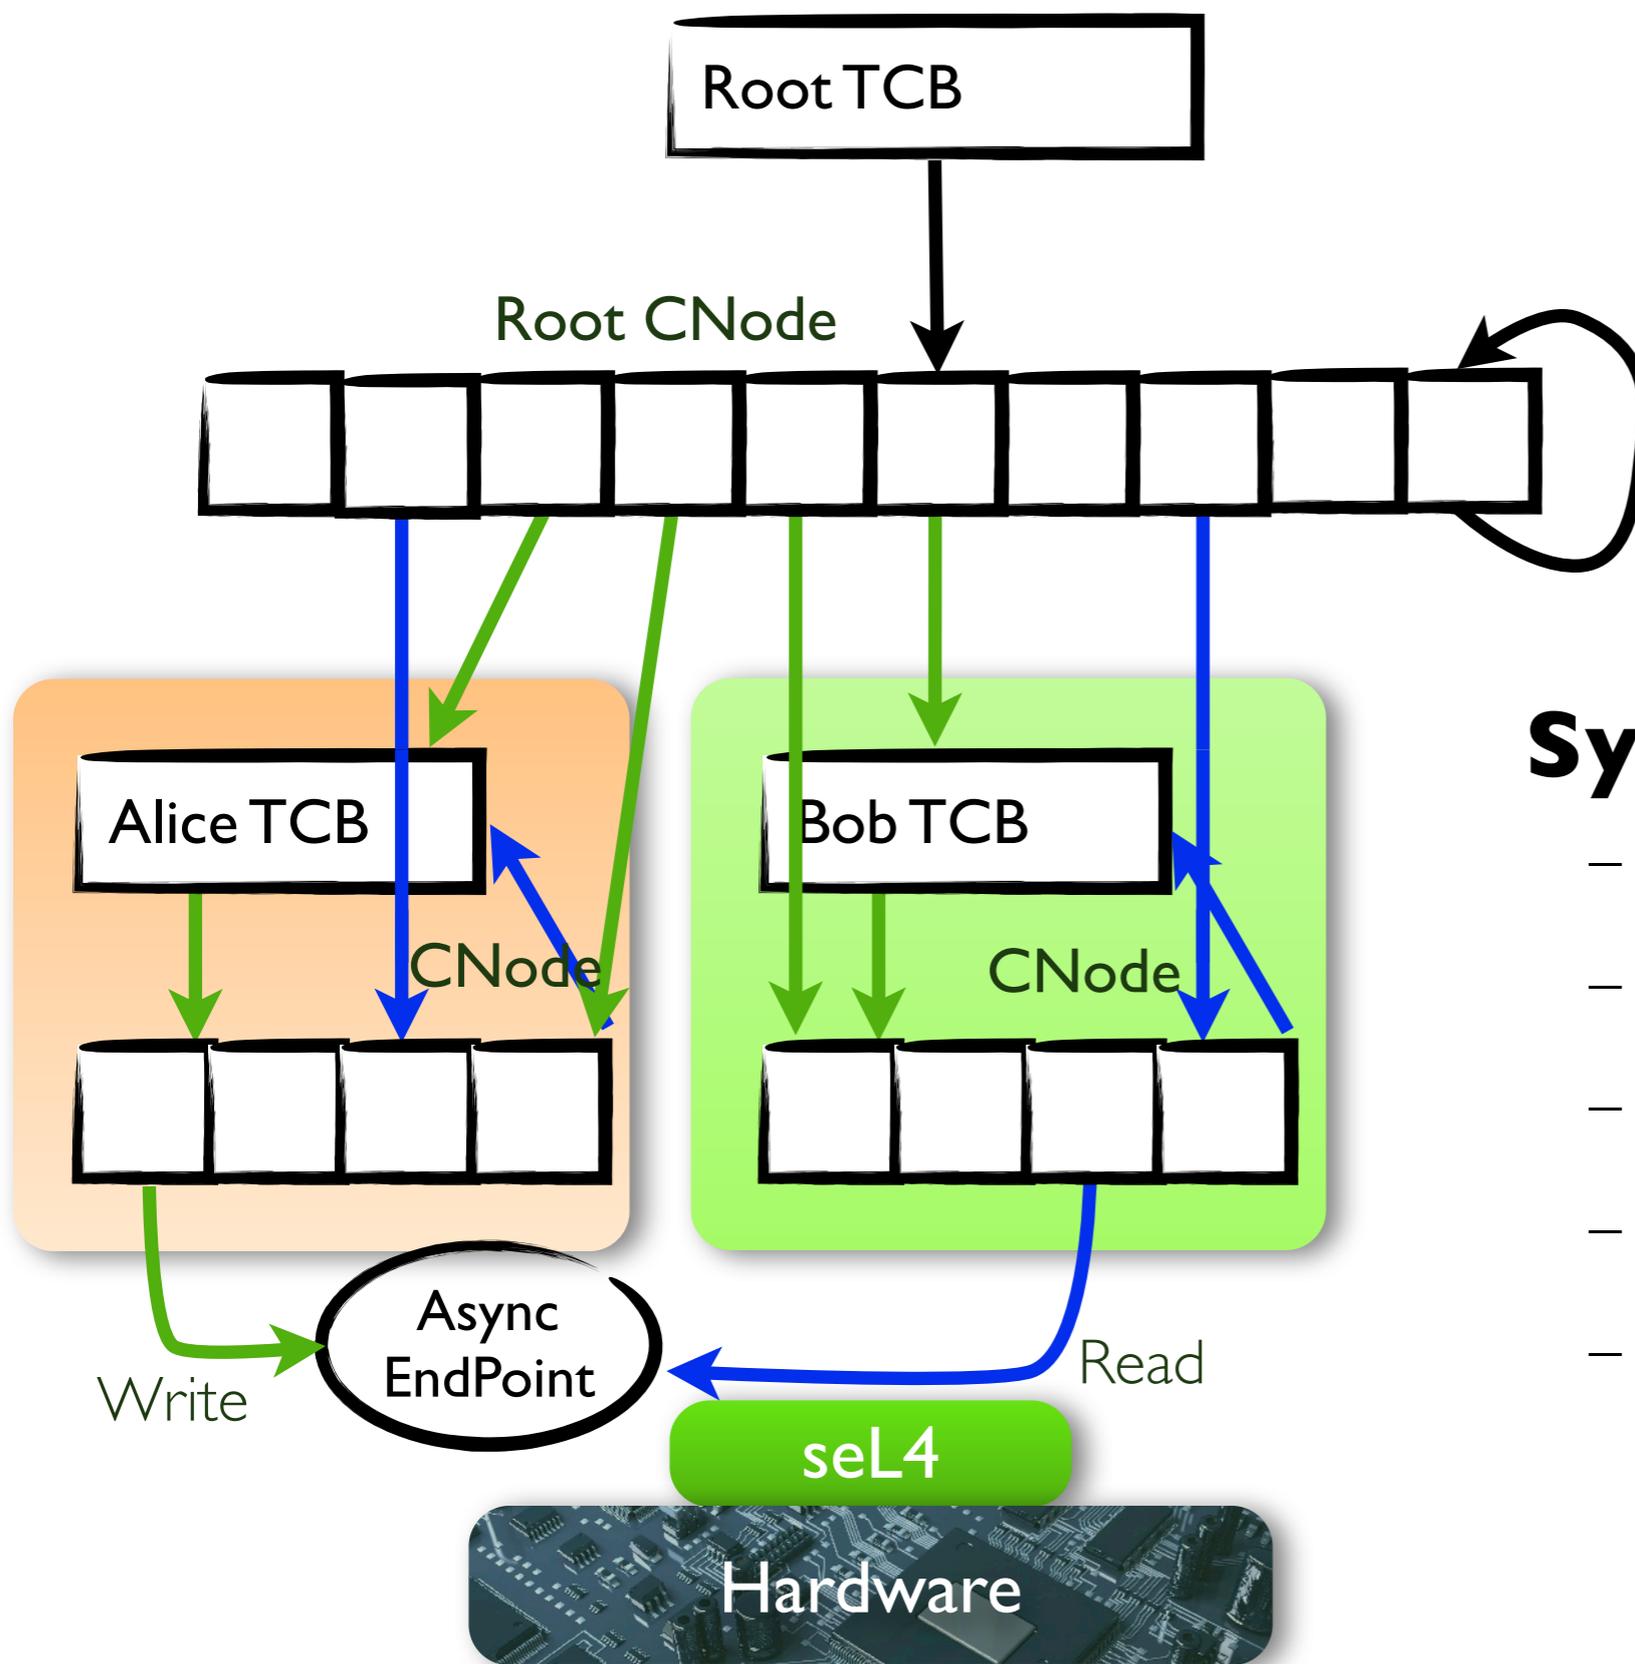
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



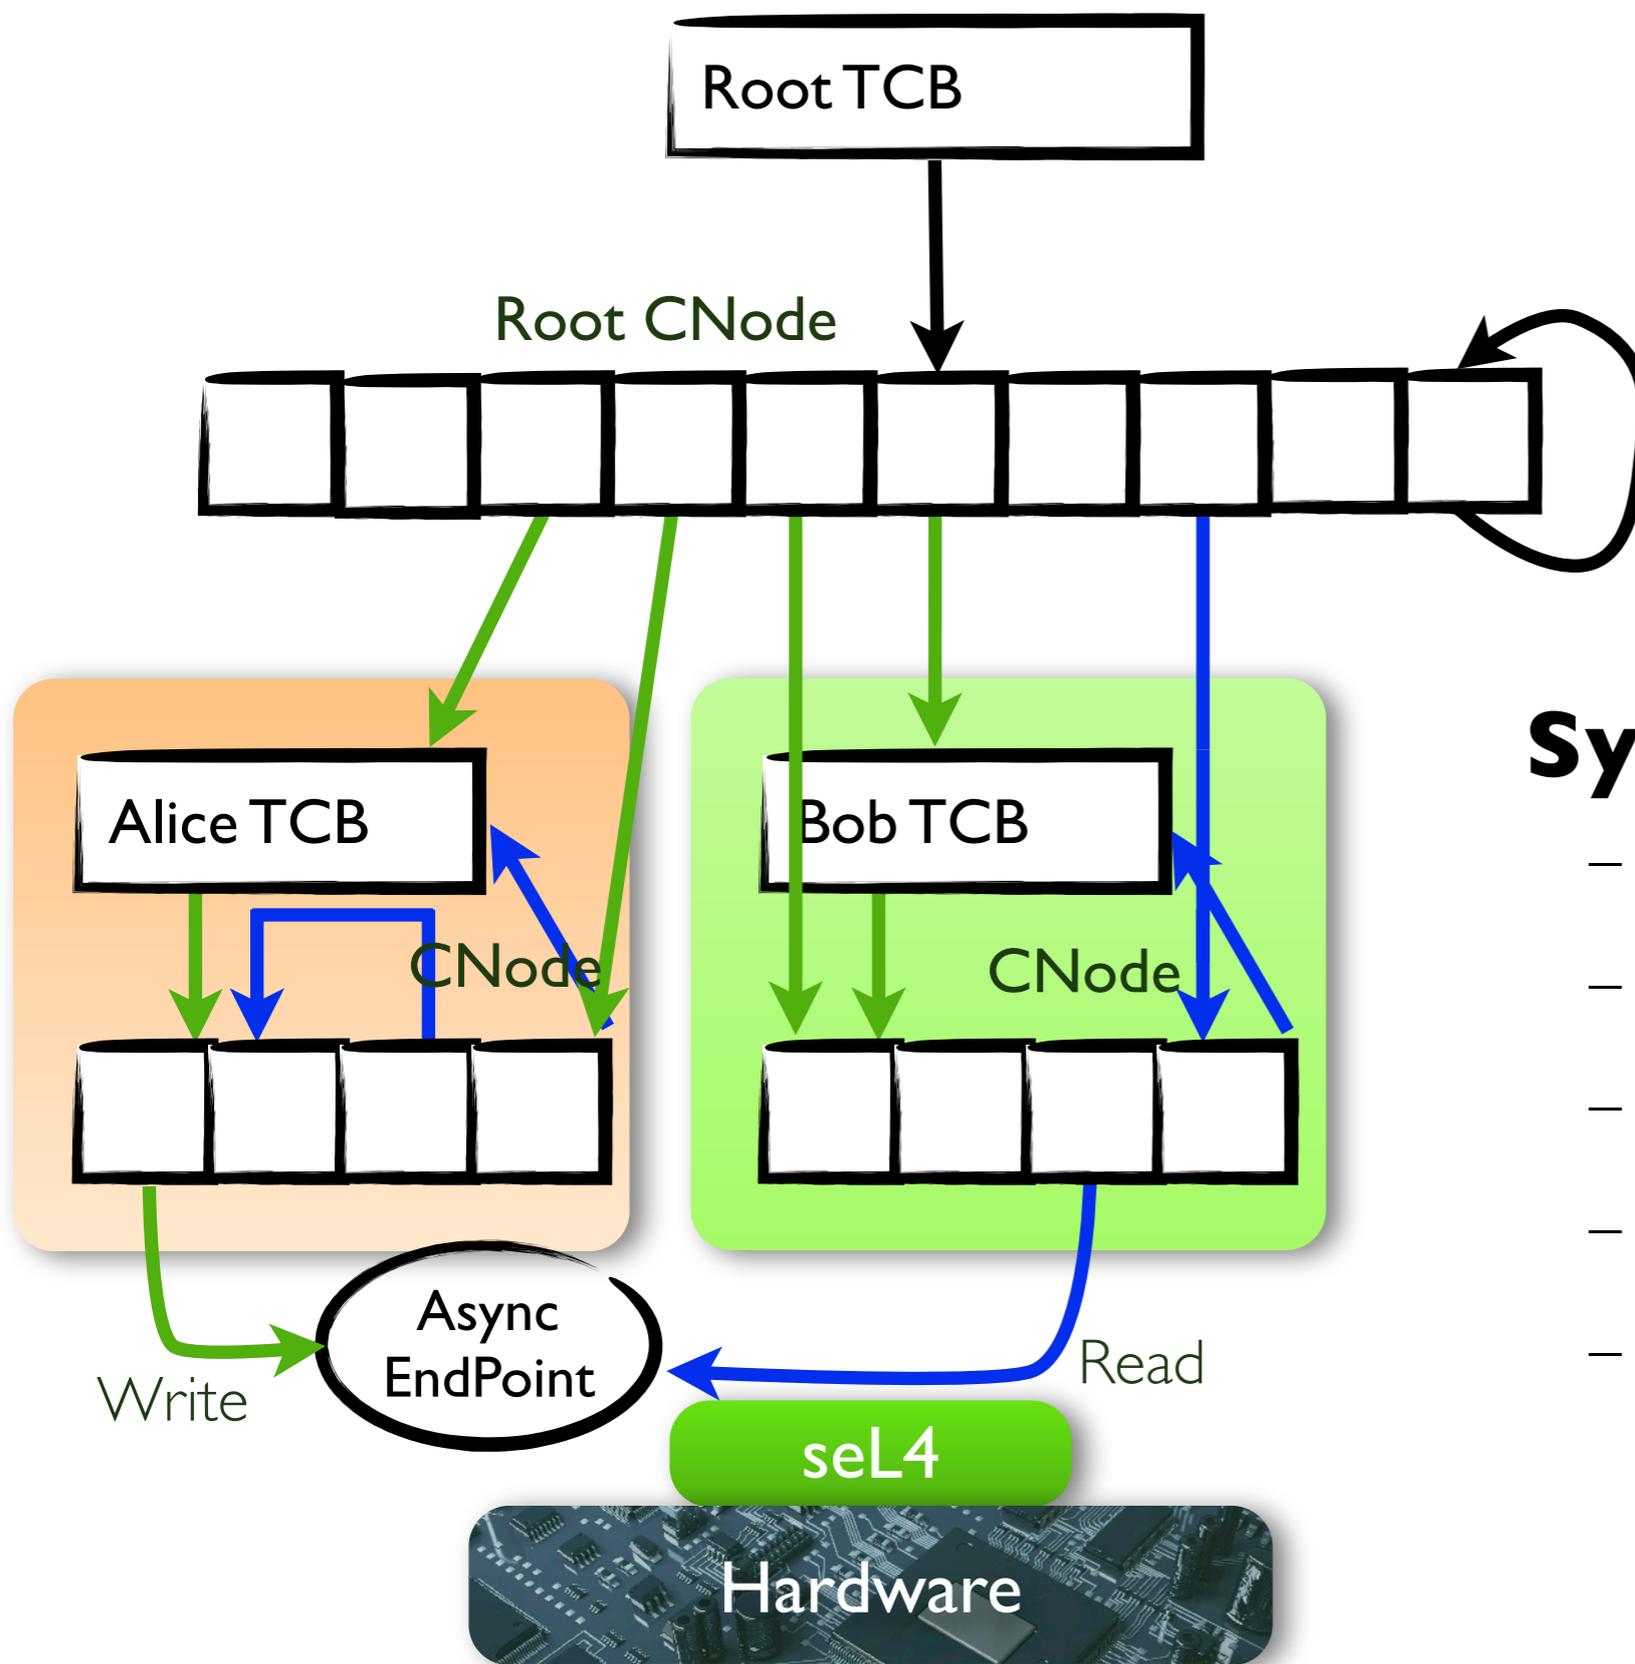
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



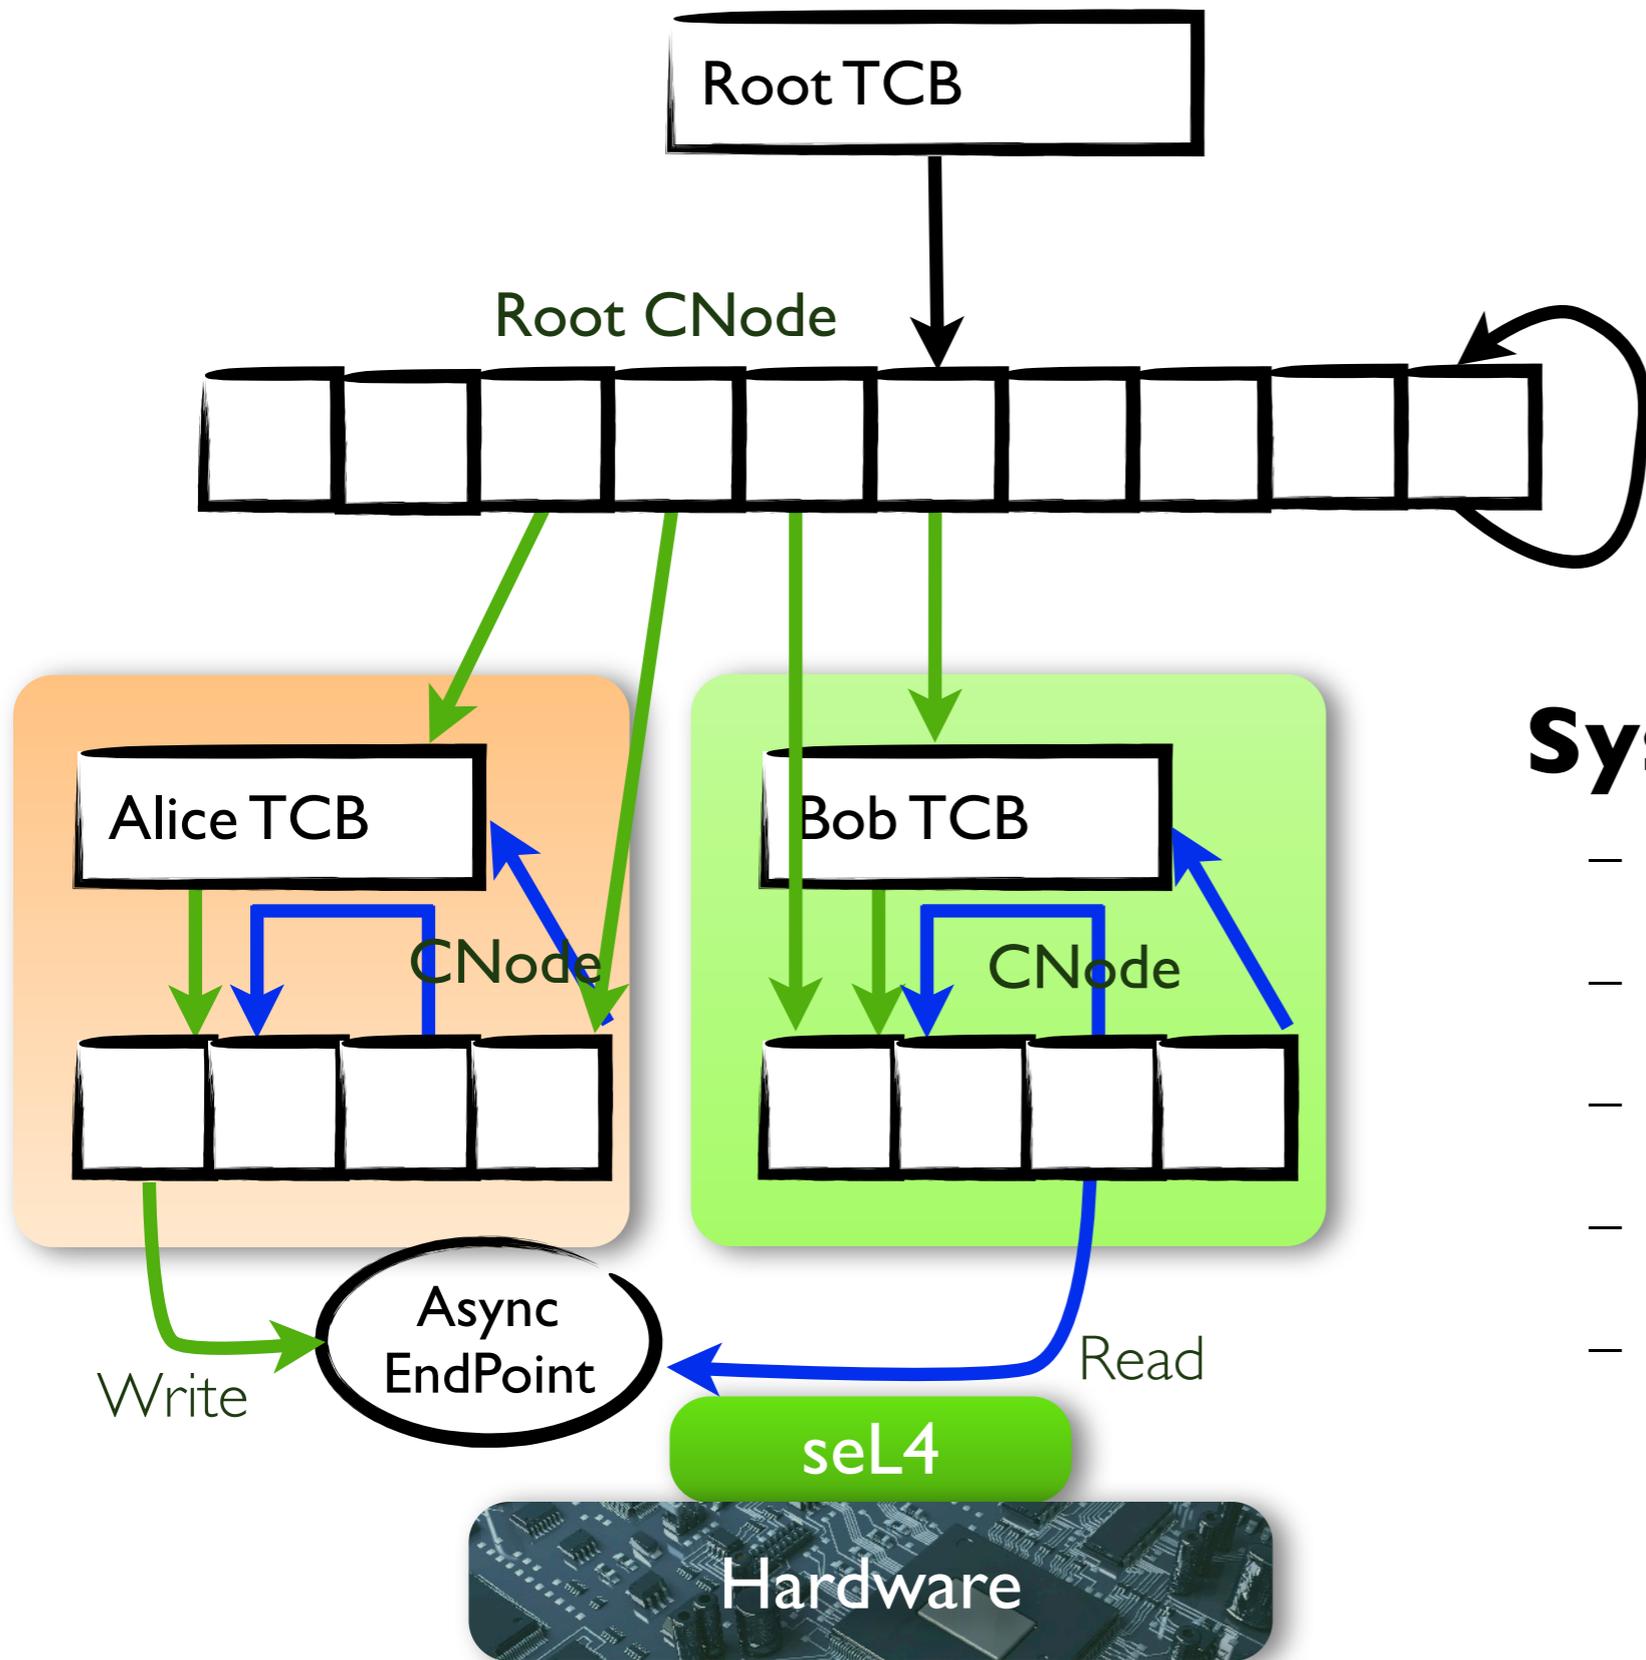
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



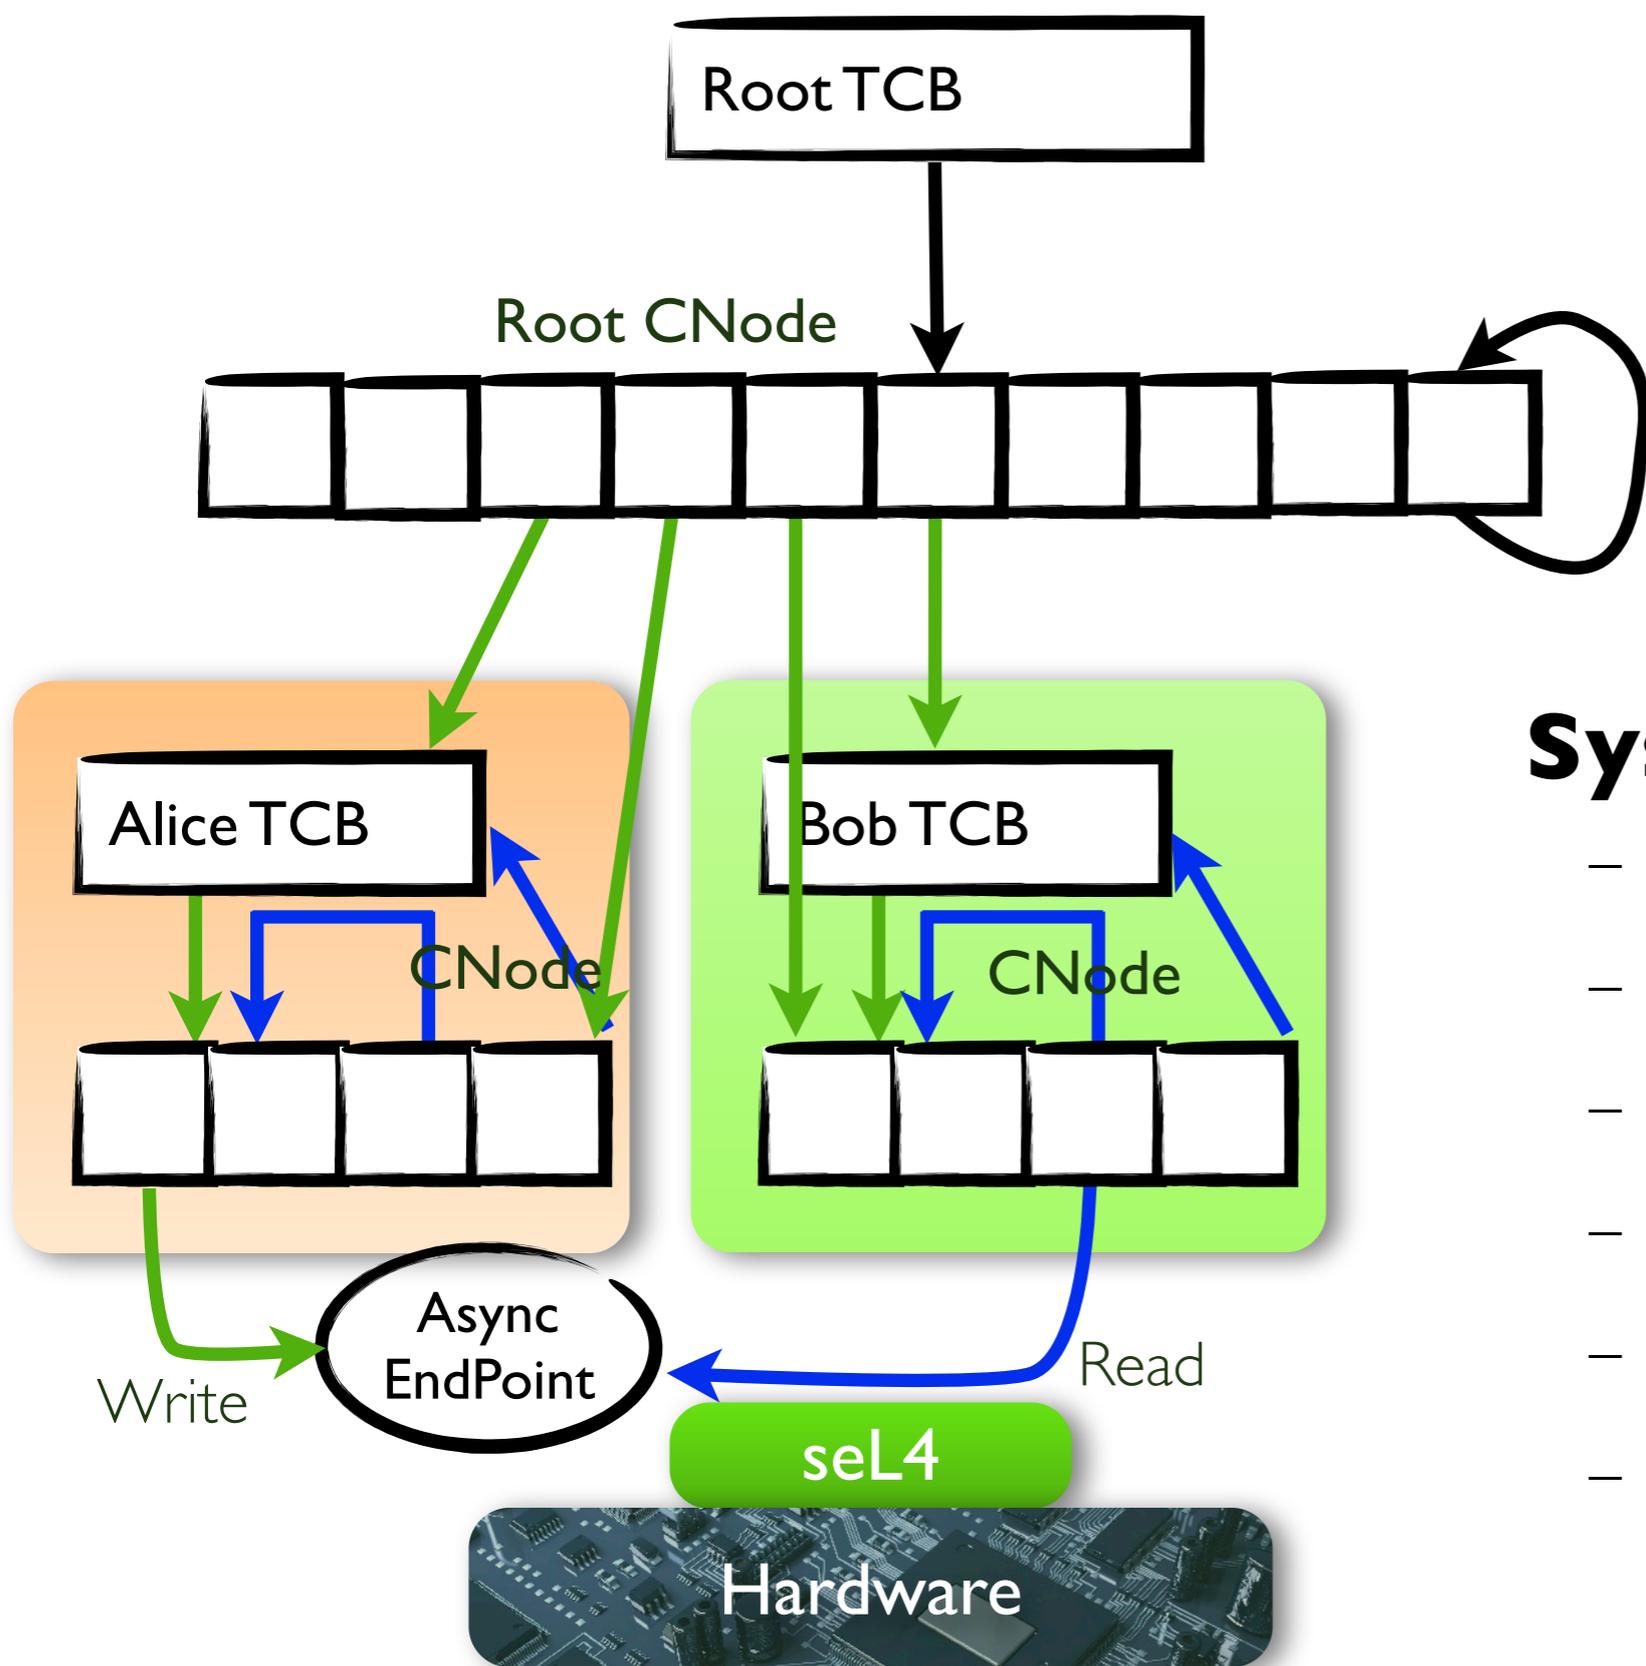
System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space



System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space
- Start threads

No objects

Empty/default objects

Duplicate capabilities to
CNodes and TCBs

Virtual memory objects
set up correctly

Virtual memory and thread
objects set up correctly

All objects set up correctly

System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space
- Start threads

Objects = Virtual Memory + TCBs + CNodes + Capless

No objects

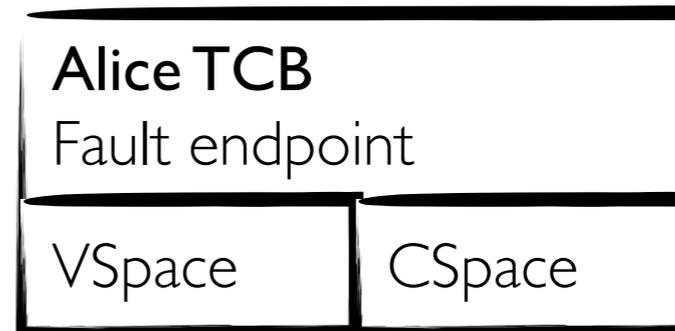
Empty/default objects

Duplicate capabilities to CNodes and TCBs

Virtual memory objects set up correctly

Virtual memory and thread objects set up correctly

All objects set up correctly



System initialisation:

- Create objects
- Duplicate capabilities
- Initialise virtual address space
- Initialise thread control blocks
- Initialise capability address space
- Start threads

Objects = Virtual Memory + TCBs + CNodes + Capless

Separation logic

- Setting up capabilities

$$\{ \text{slot}_1 \mapsto_c \text{cap} \wedge^* \text{slot}_2 \mapsto_c - \wedge^* R \}$$

move_cap slot₁ slot₂

$$\{ \text{slot}_1 \mapsto_c \text{NullCap} \wedge^* \text{slot}_2 \mapsto_c \text{cap} \wedge^* R \}$$

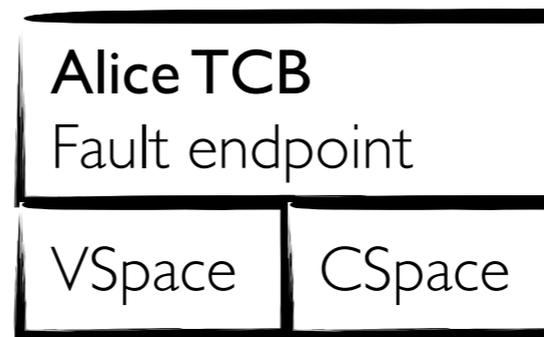
- Monadic maps (loops)

$$\left(\wedge R \ x \in xs. \{ P \ x \ \wedge^* R \} \ f \ x \ \{ Q \ x \ \wedge^* R \} \right) \implies$$

$$\{ \wedge^* \text{map } P \ xs \ \wedge^* R \} \ \text{mapM } f \ xs \ \{ \wedge^* \text{map } Q \ xs \ \wedge^* R \}$$

Separation logic

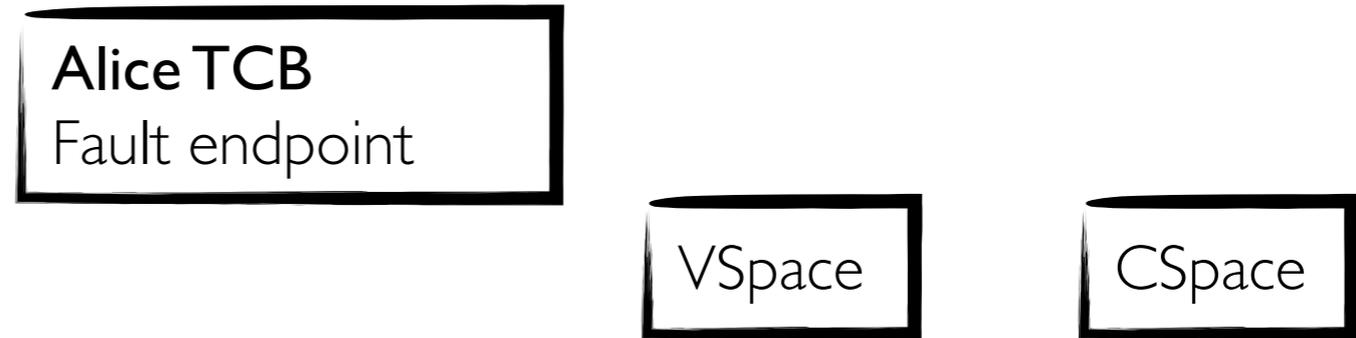
CapDL object heap:
cdl_object_id \Rightarrow cdl_object option



0x20

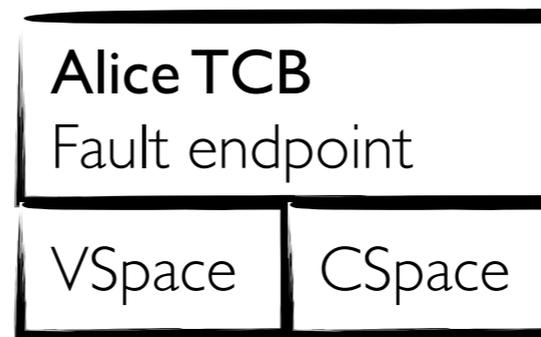
Separation logic

Separation heap:



CapDL object heap:

`cdl_object_id` \Rightarrow `cdl_object` option



0x20



Separation logic

Separation heap:

(cdl_object_id × cdl_component)
⇒ sep_entity option



(0x20, Fields)



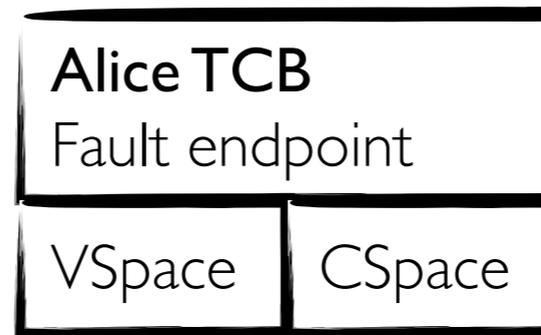
(0x20, Slot 0)



(0x20, Slot 1)

CapDL object heap:

cdl_object_id ⇒ cdl_object option



0x20



Separation logic

Separation heap:

(cdl_object_id × cdl_component)
⇒ sep_entity option



(0x20, Fields)



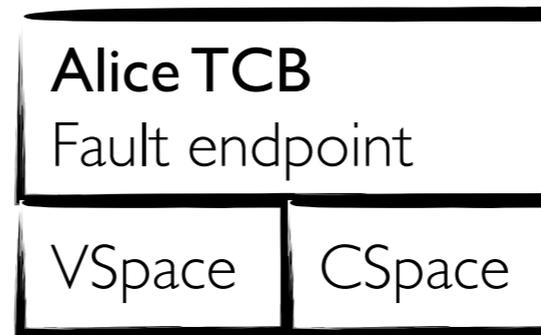
(0x20, Slot 0)



(0x20, Slot 1)

CapDL object heap:

cdl_object_id ⇒ cdl_object option



0x20



Instantiate separation heap in the standard manner

Overall theorem

If well_formed spec **and**

obj_ids = dom (cdl_objects spec) **and**

distinct obj_ids **then**

{ «valid_boot_info bootinfo spec » }

init_system spec bootinfo obj_ids

{ $\lambda s. \exists \varphi. \ll \wedge^* \text{map} (\text{object_initialised spec } \varphi) \text{ obj_ids } \wedge^*$

$\text{si_objects spec } \varphi \gg s \wedge$

$\text{injective } \varphi \wedge \text{dom } \varphi = \text{obj_ids} \}$

Proof decomposition



$\llbracket \text{dom (slots_of obj_id spec) = slots; distinct slots} \rrbracket$

$\implies \text{object_initialised spec } \varphi \text{ obj_id} =$

$(\text{object_fields_initialised spec } \varphi \text{ obj_id} \wedge^*$

$\wedge^* \text{map (object_slot_initialised spec } \varphi \text{ obj_id) slots} \wedge^*$

$\text{object_empty_slots_initialised spec } \varphi \text{ obj_id})$

Proof decomposition

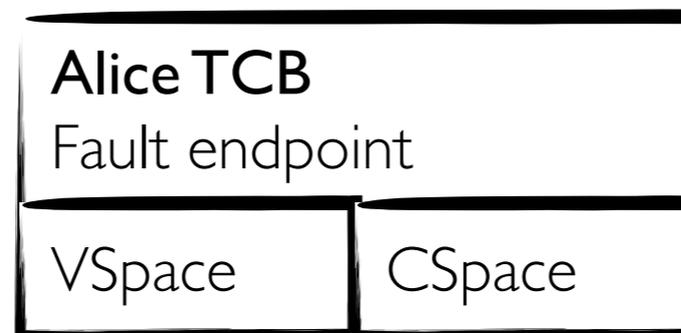
$\llbracket \text{dom (slots_of obj_id spec) = slots; distinct slots} \rrbracket$

$\implies \text{object_initialised spec } \varphi \text{ obj_id} =$

$(\text{object_fields_initialised spec } \varphi \text{ obj_id} \wedge^*$

$\wedge^* \text{map (object_slot_initialised spec } \varphi \text{ obj_id) slots} \wedge^*$

$\text{object_empty_slots_initialised spec } \varphi \text{ obj_id})$



Proof decomposition

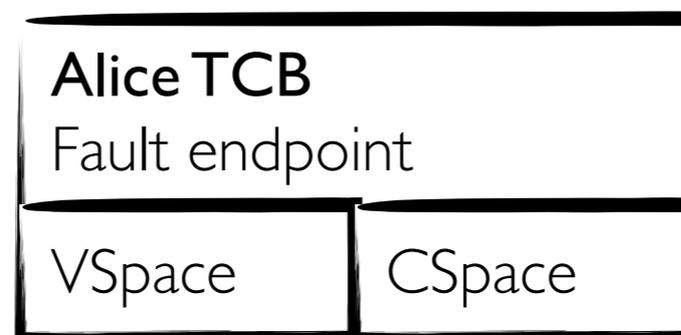
$\llbracket \text{dom (slots_of obj_id spec) = slots; distinct slots} \rrbracket$

$\Rightarrow \text{object_initialised spec } \varphi \text{ obj_id} =$

$(\text{object_fields_initialised spec } \varphi \text{ obj_id} \wedge^*$

$\wedge^* \text{map (object_slot_initialised spec } \varphi \text{ obj_id) slots} \wedge^*$

$\text{object_empty_slots_initialised spec } \varphi \text{ obj_id})$



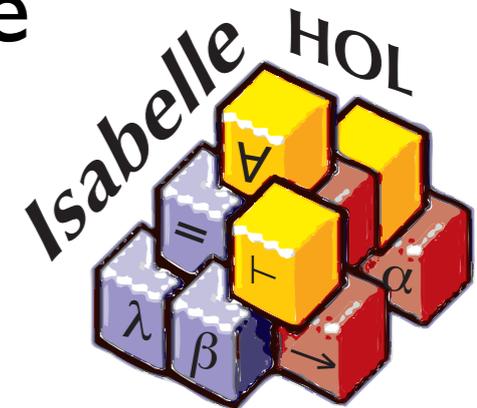
Connect to seL4 kernel API specifications we developed

Future work

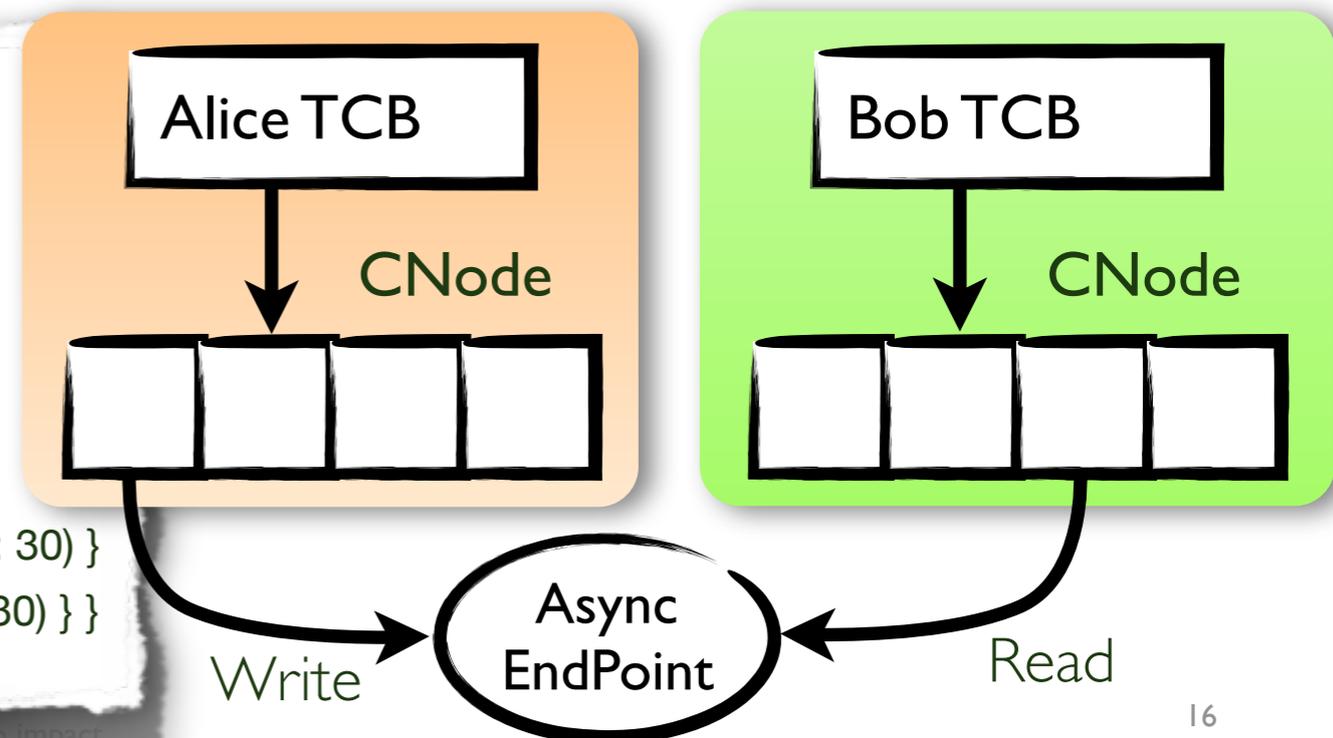
- **Assumptions/limitations**
 - Scheduling
 - Don't support all capability types
- **Refinement proof of C implementation**

Summary

- Verified formal algorithm for automatically initialising capability based systems with high-assurance
- Proven in Isabelle/HOL
 - System initialiser specification ~ 400 lines
 - Kernel API specification proofs ~ 10,500 lines
 - Proof of system initialiser specification ~ 8,100 lines
- Linked capDL descriptions to access control policy



```
objects { cnode_alice = cnode (2 bits)
          tcb_alice = tcb ()
          cnode_bob = cnode (2 bits)
          tcb_bob = tcb ()
          aep_shared = aep }
caps { cnode_alice { 0x0: aep_shared (W) }
       cnode_bob { 0x2: aep_shared (R) }
       tcb_alice { cspace: cnode_alice (guard: 0, guard_size: 30) }
       tcb_bob { cspace: cnode_bob (guard: 0, guard_size: 30) } }
```



Questions?

Formally Verified System Initialisation

Andrew Boyton

June Andronick, Callum Bannister,
Matthew Fernandez, Xin Gao,
David Greenaway, Gerwin Klein,
Corey Lewis and Thomas Sewell