



Let's build secure systems on a correct kernel

June Andronick



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

NICTA Funding and Supporting Members and Partners



Credits



Gerwin Klein

(part of) the Trustworthy Embedded Systems crowd



Secure Systems "à la NICTA"

For 1 critical system:

- 1 desired security property P
- an interactive theorem prover
- a bit of patience

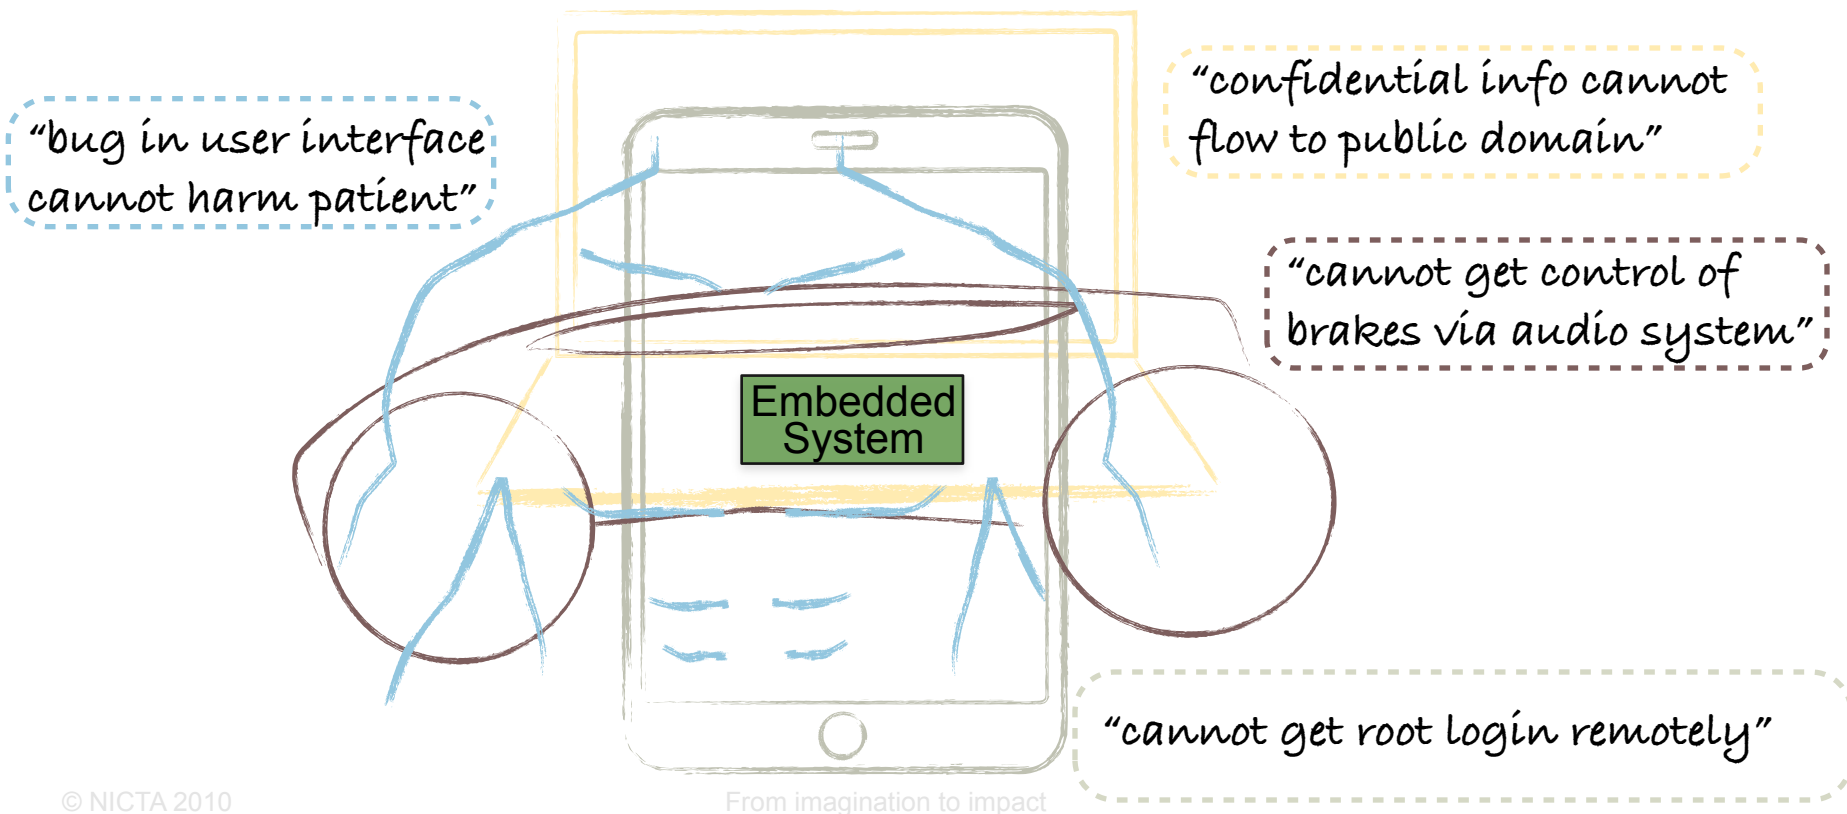
1. carefully design your system
2. prove that the design enforces P
3. prove correctness of the TCB
4. prove isolation

Motivation


Aim: Trustworthy Embedded Systems

Target: code-level guarantee
at reasonable cost

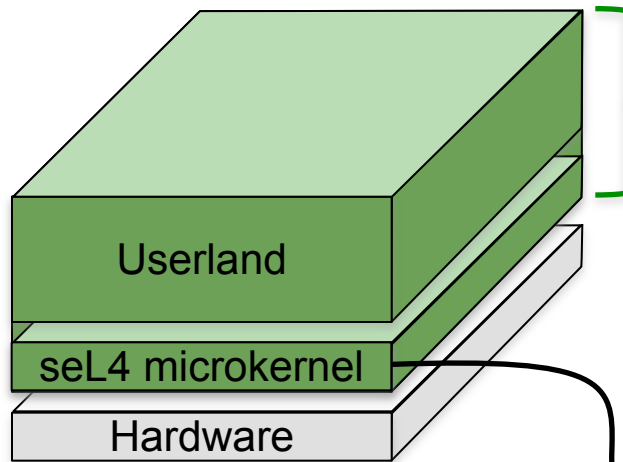
real-world, usable systems
(low-level, performant)



Approach



Approach



How to prove this is trustworthy?

1. Trustworthy foundation → **seL4**

functional correctness for 10,000 loc

Formal functional spec

Formal proof of refinement

Code



4.Verified

Result: “Every behavior of the code is a behavior of the spec”


Corollary: “execution always defined”
(no buffer overflows, ...)

Assumptions:

- compiler + linker
- assembly code (600 loc)
- hardware (ARMv6)
- cache/TLB
- boot code (1,200 loc)

WORK IN
PROGRESS

Approach



How to prove this is trustworthy?

1. Trustworthy foundation → **seL4**

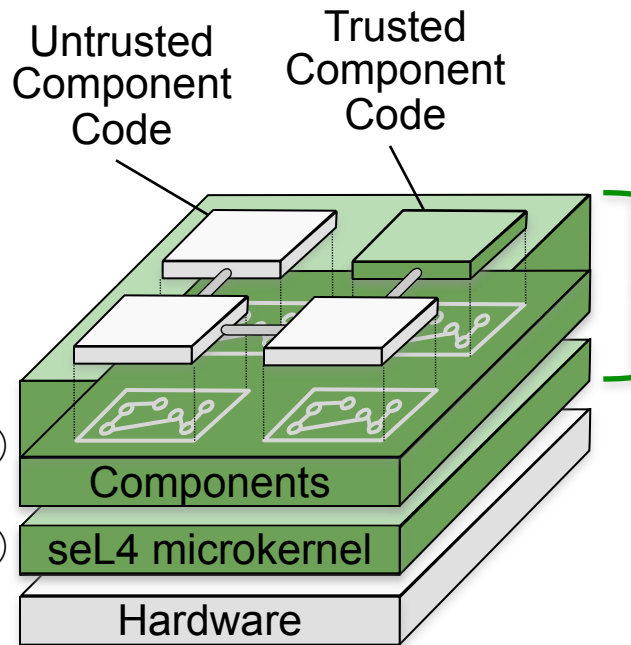
functional correctness for 10,000 loc

2. Strategic componentized security architecture

formal guarantees for >1,000,000 loc

Idea: Strong *guarantees* about *whole system*
without needing to reason about all of its code

Approach



How to prove this is trustworthy?


1. Trustworthy foundation → **seL4**
functional correctness for 10,000 loc

2. Strategic componentized security architecture
formal guarantees for >1,000,000 loc

Idea: Strong *guarantees* about *whole system*
without needing to reason about all of its code

How: Using seL4's access control (*capabilities*)

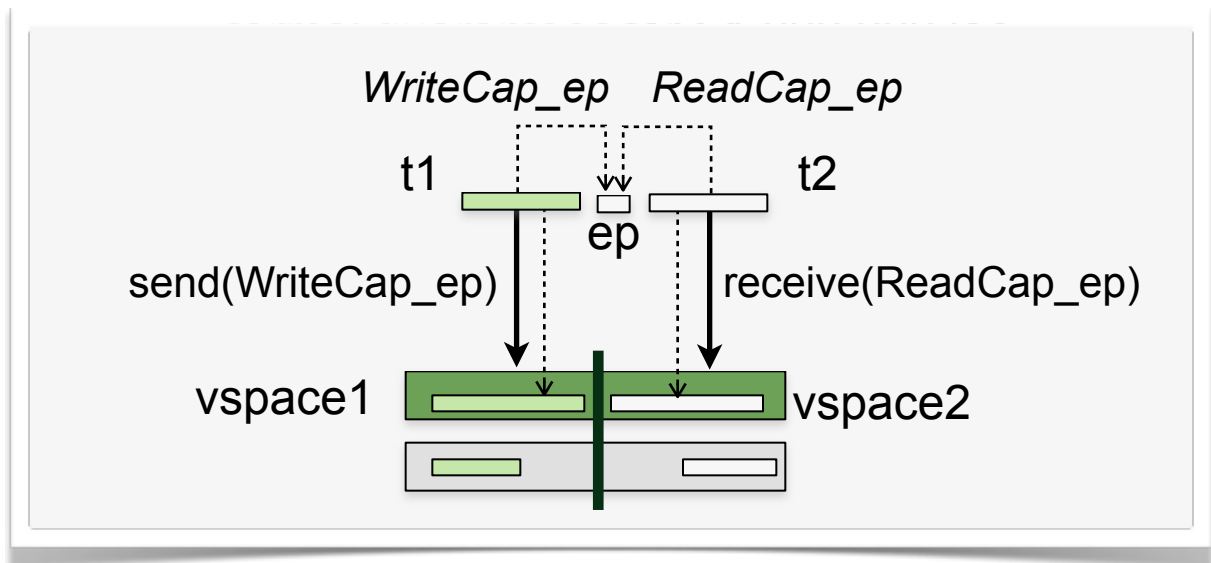
Approach



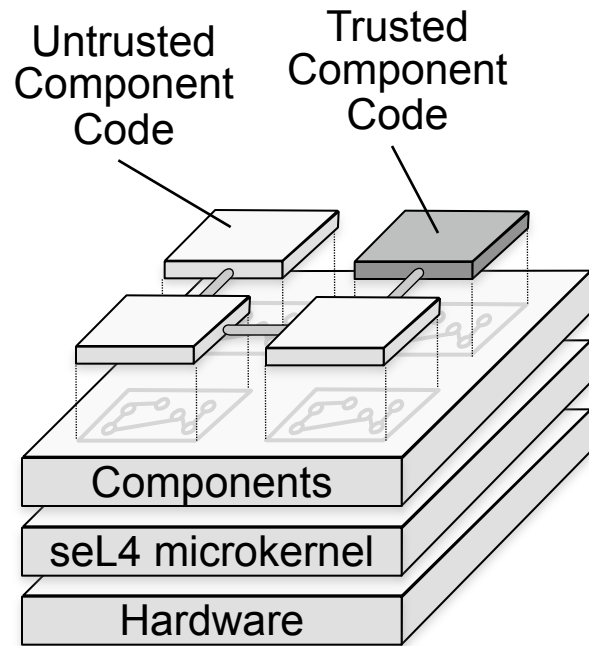
How to prove this is trustworthy?

1. Trustworthy foundation → **seL4**
functional correctness for 10,000 loc

2. Strategic componentized security architecture




Careful design



System Implementation

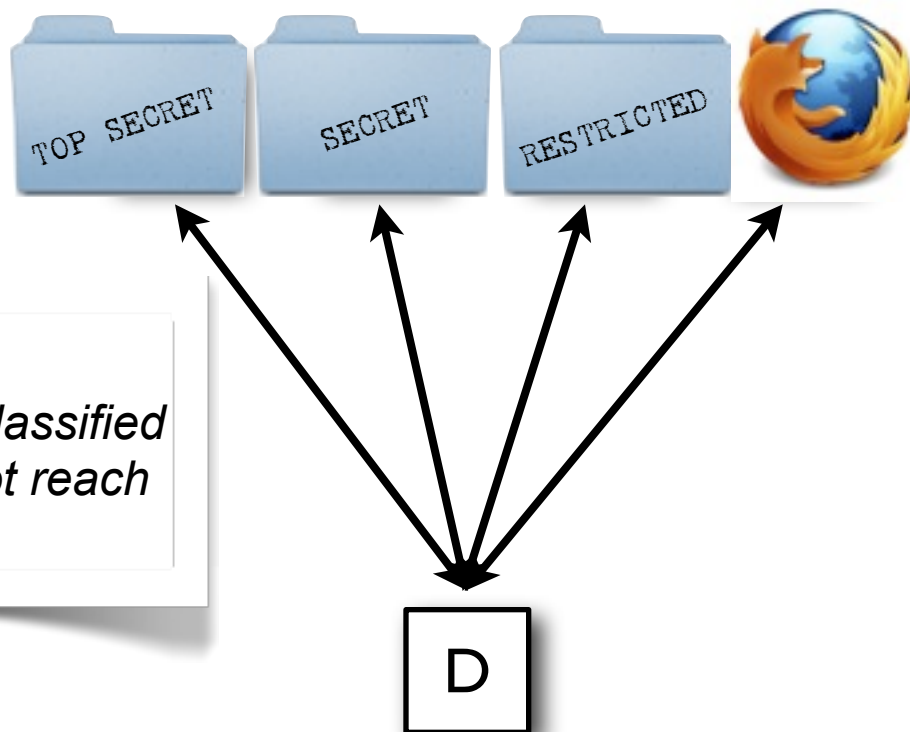
Security Architecture



1a. minimal Trusted Computing Base

Case-study

Classified Networks



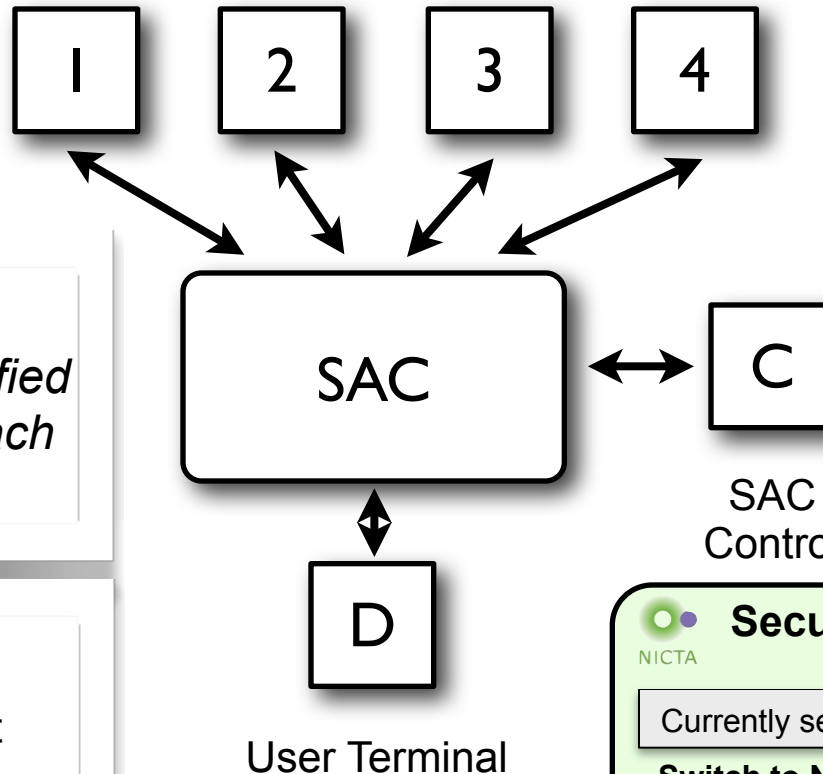
Goal:

Data from one classified network must not reach another

User Terminal

Secure Access Controller (SAC)

Classified Networks

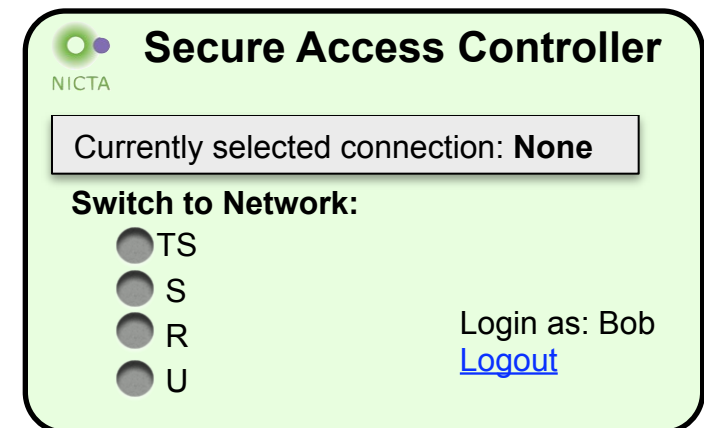


Goal:

Data from one classified network must not reach another

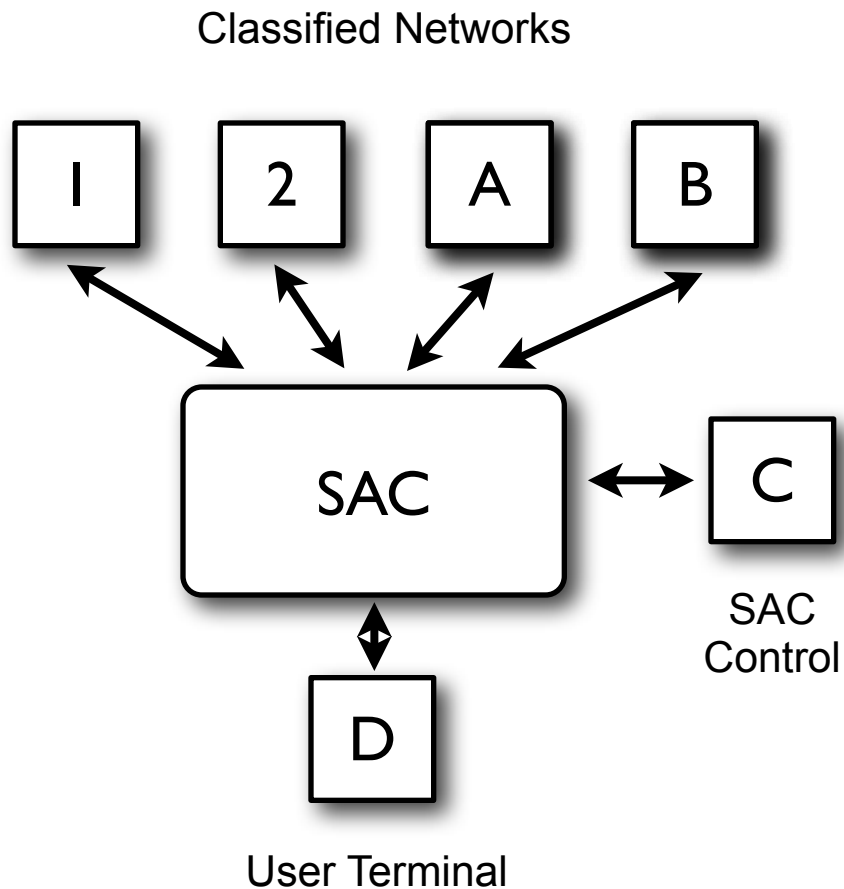
Assumptions:

- User terminal will not leak data
- All networks are otherwise malicious

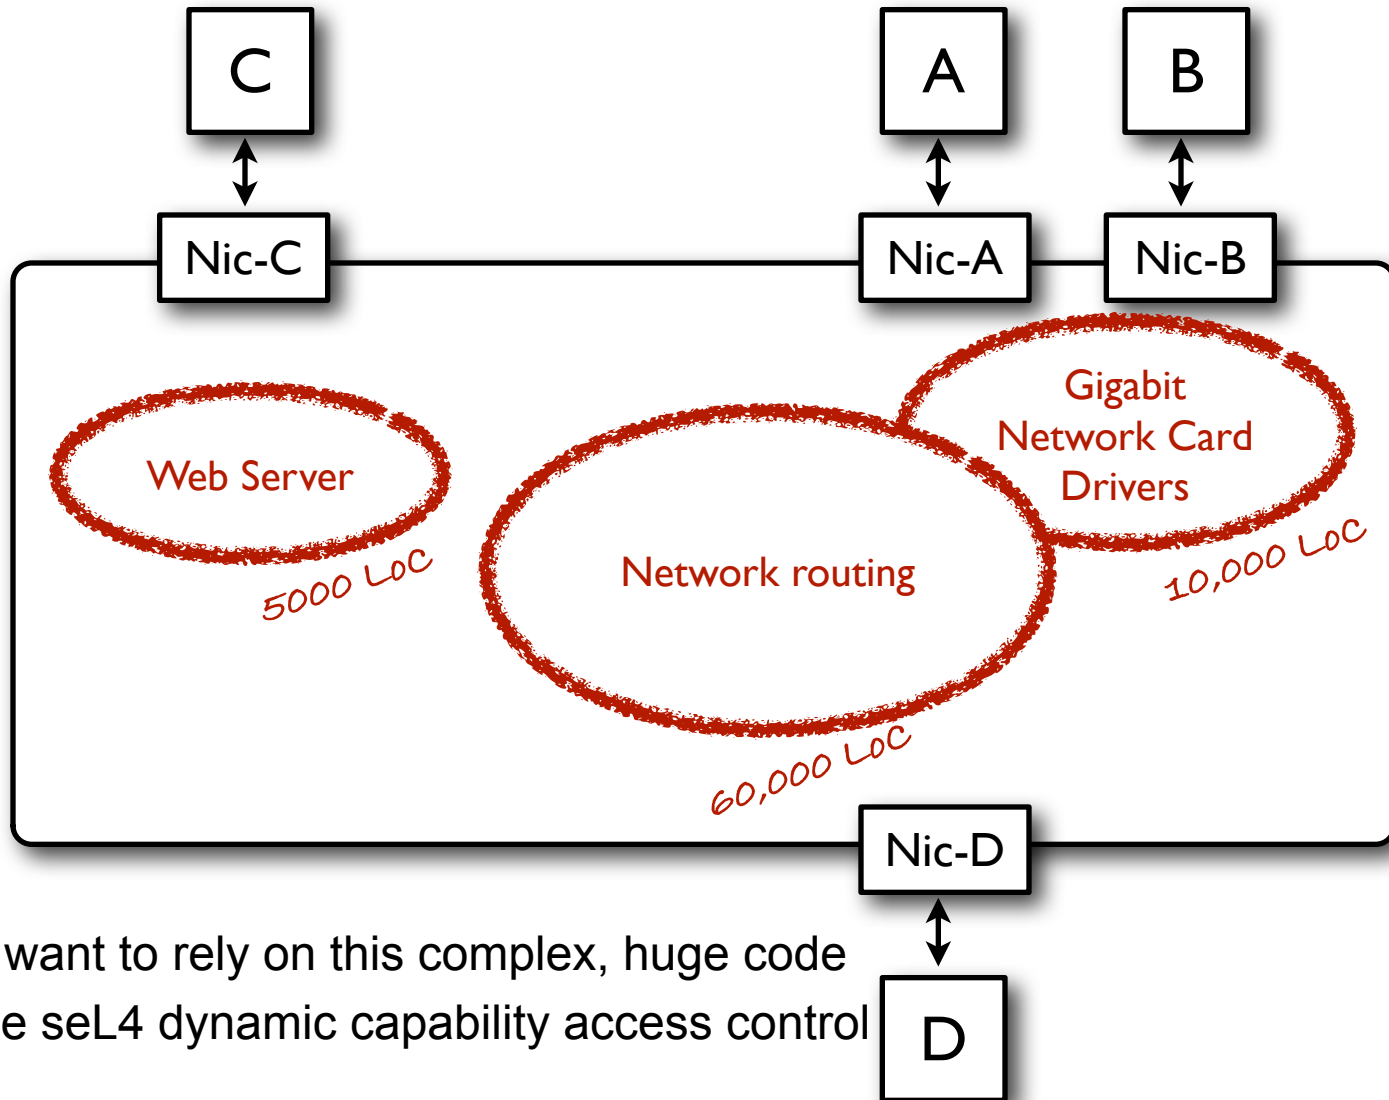


The screenshot shows the user interface of the Secure Access Controller. It features the NICTA logo and the title 'Secure Access Controller'. Below the title, it displays 'Currently selected connection: None'. Under the heading 'Switch to Network:', there are four radio button options: 'TS', 'S', 'R', and 'U'. On the right side, it shows 'Login as: Bob' and a blue 'Logout' link.

Design

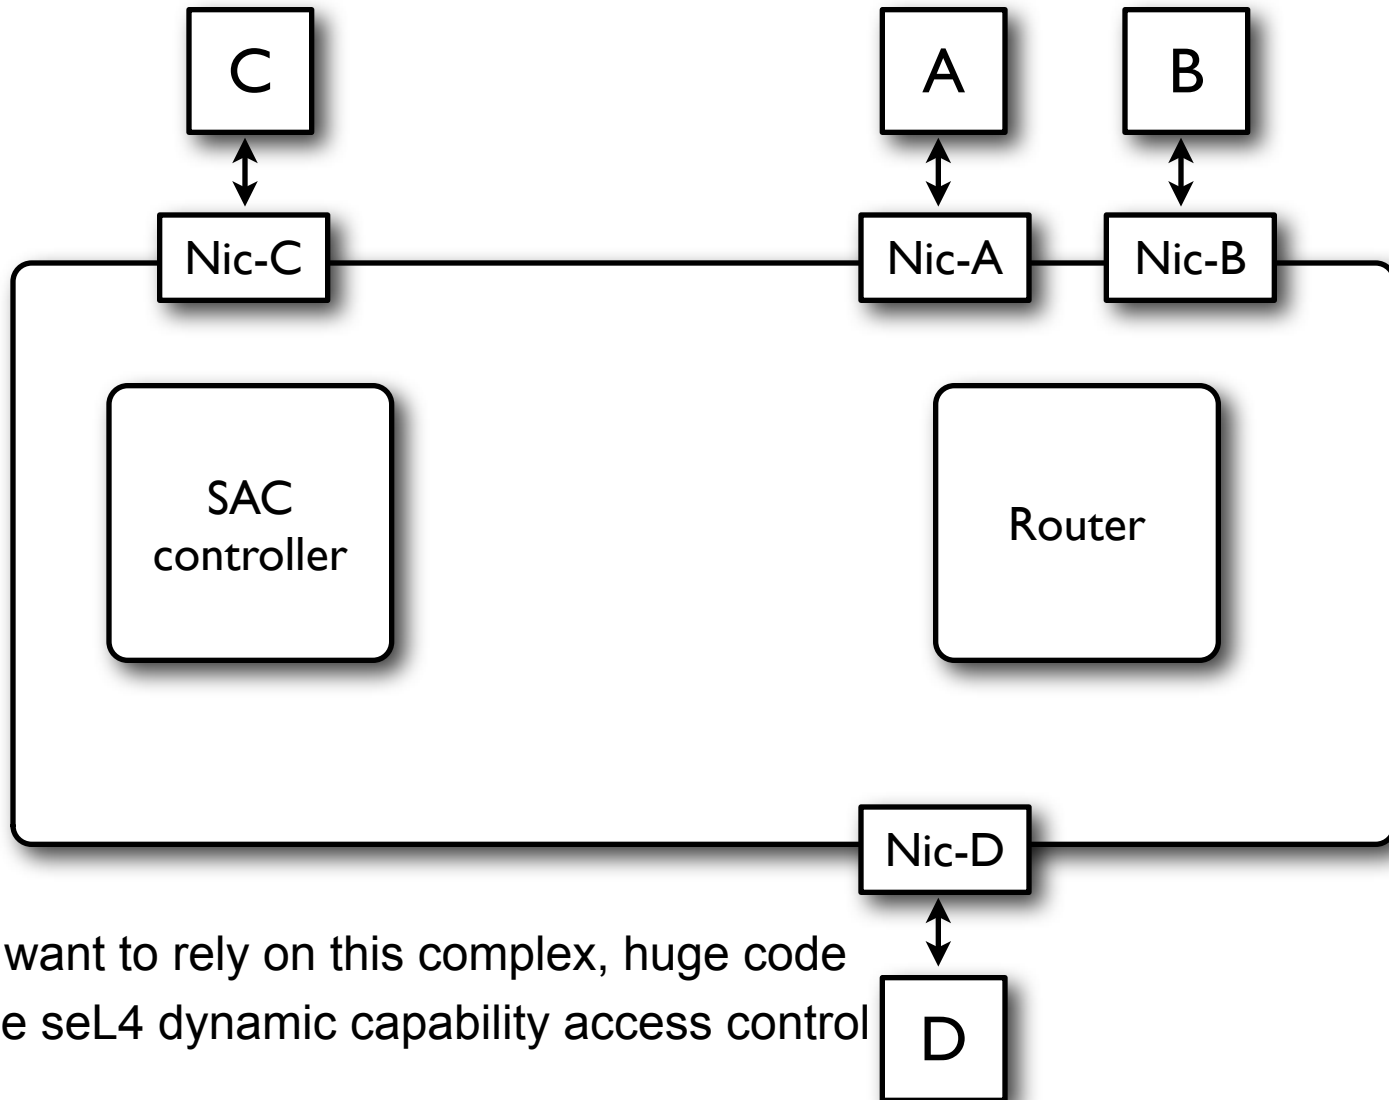


Minimal TCB



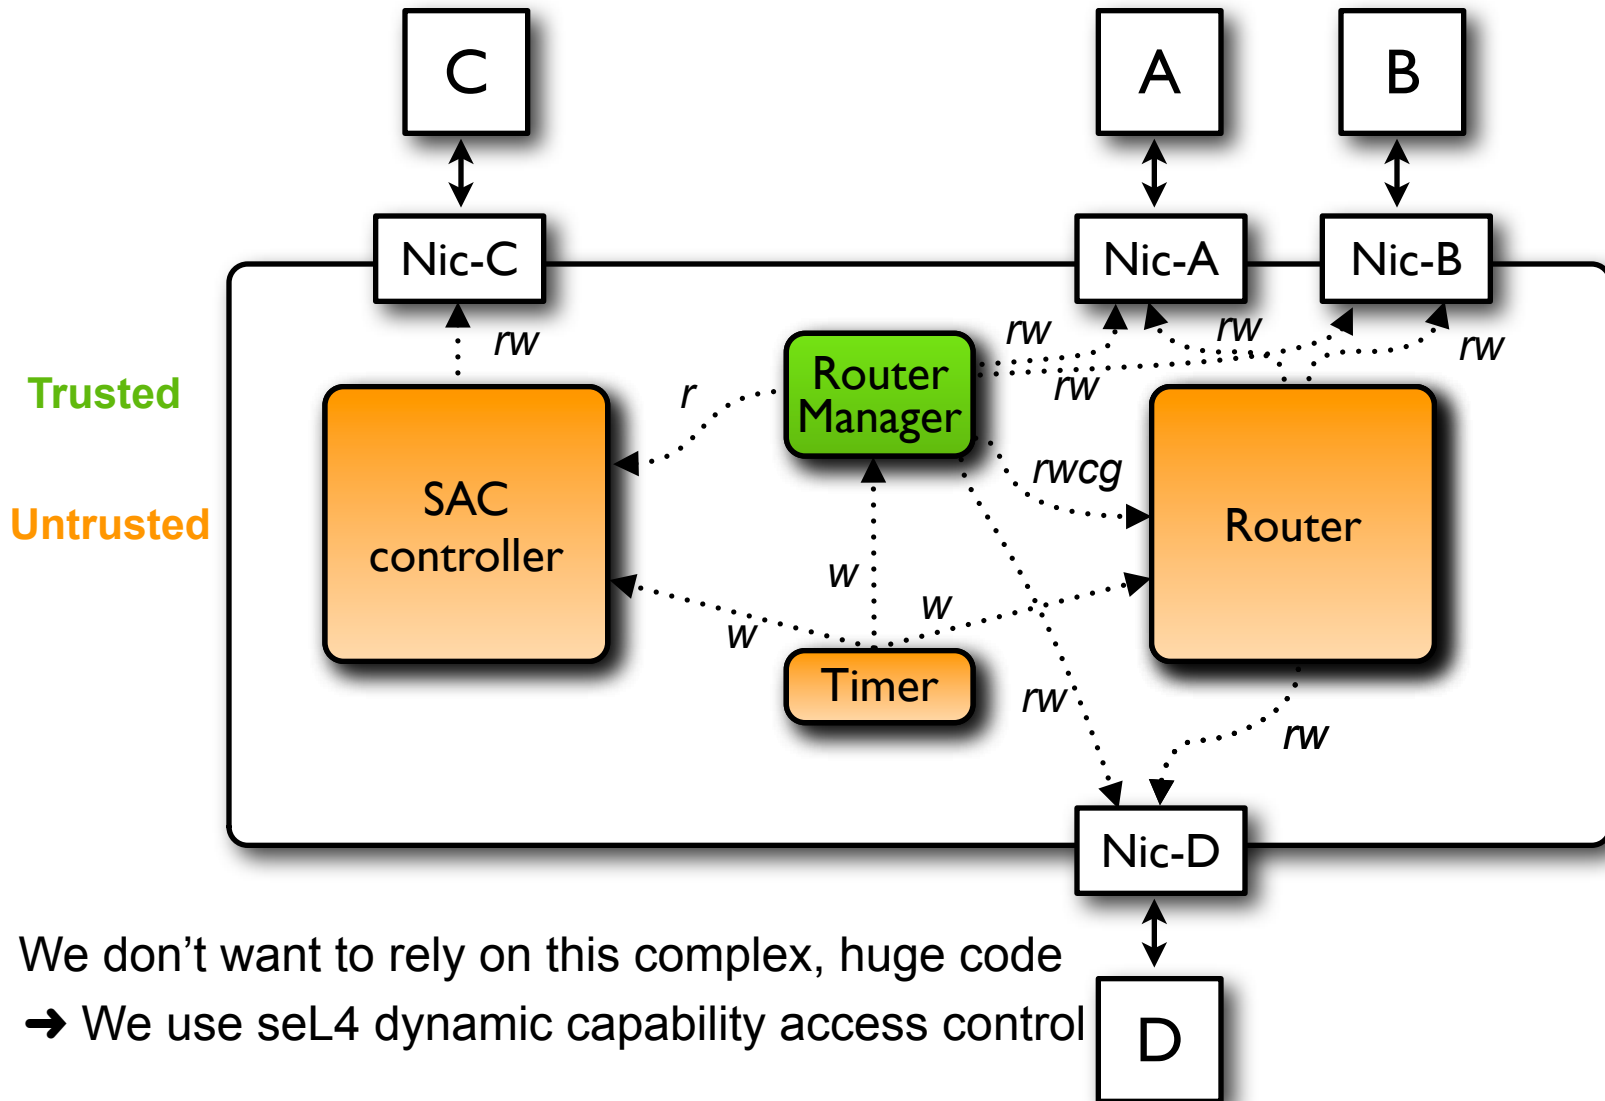
We don't want to rely on this complex, huge code
→ We use seL4 dynamic capability access control

Minimal TCB



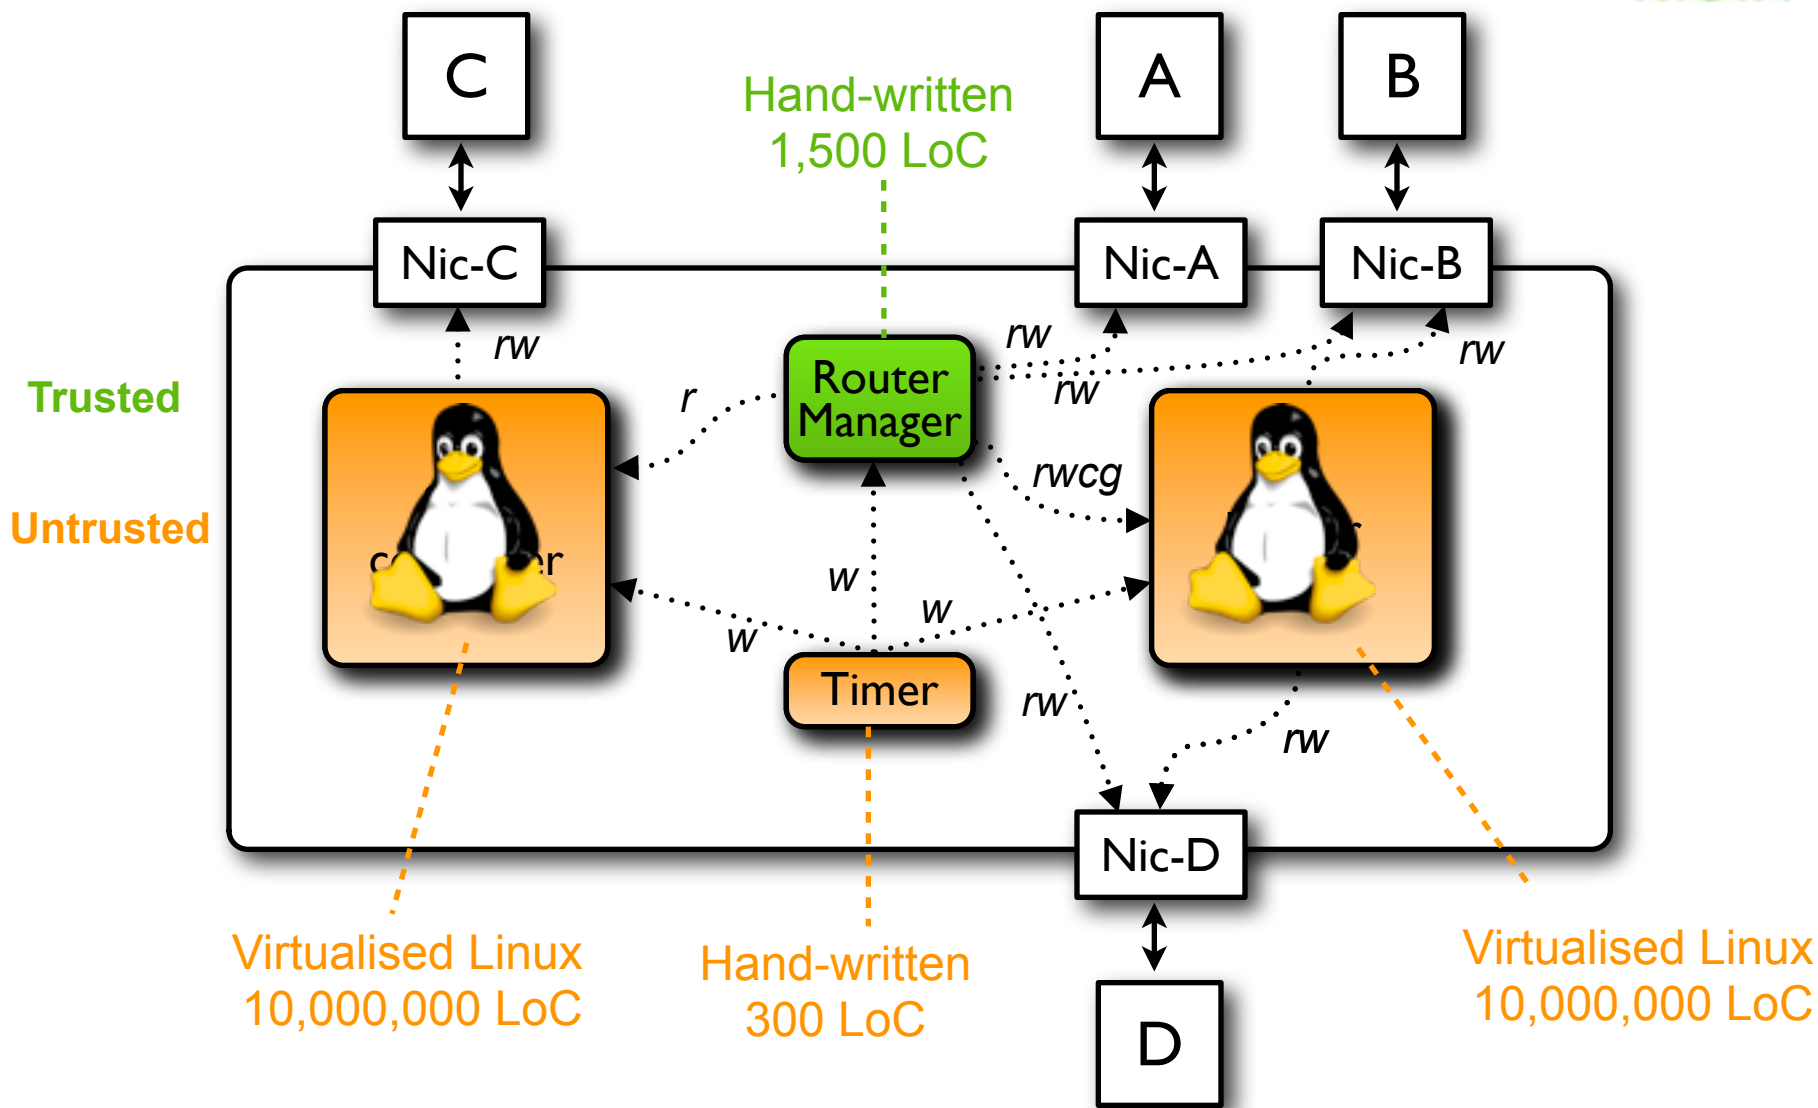
We don't want to rely on this complex, huge code
→ We use seL4 dynamic capability access control

Minimal TCB


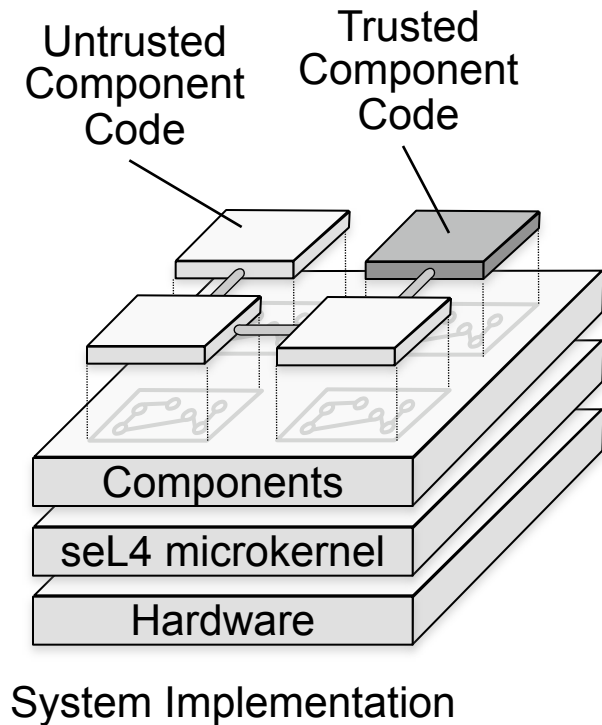


We don't want to rely on this complex, huge code
→ We use seL4 dynamic capability access control

Minimal TCB: Implementation



Back to the general picture



1a. minimal Trusted Computing Base ✓

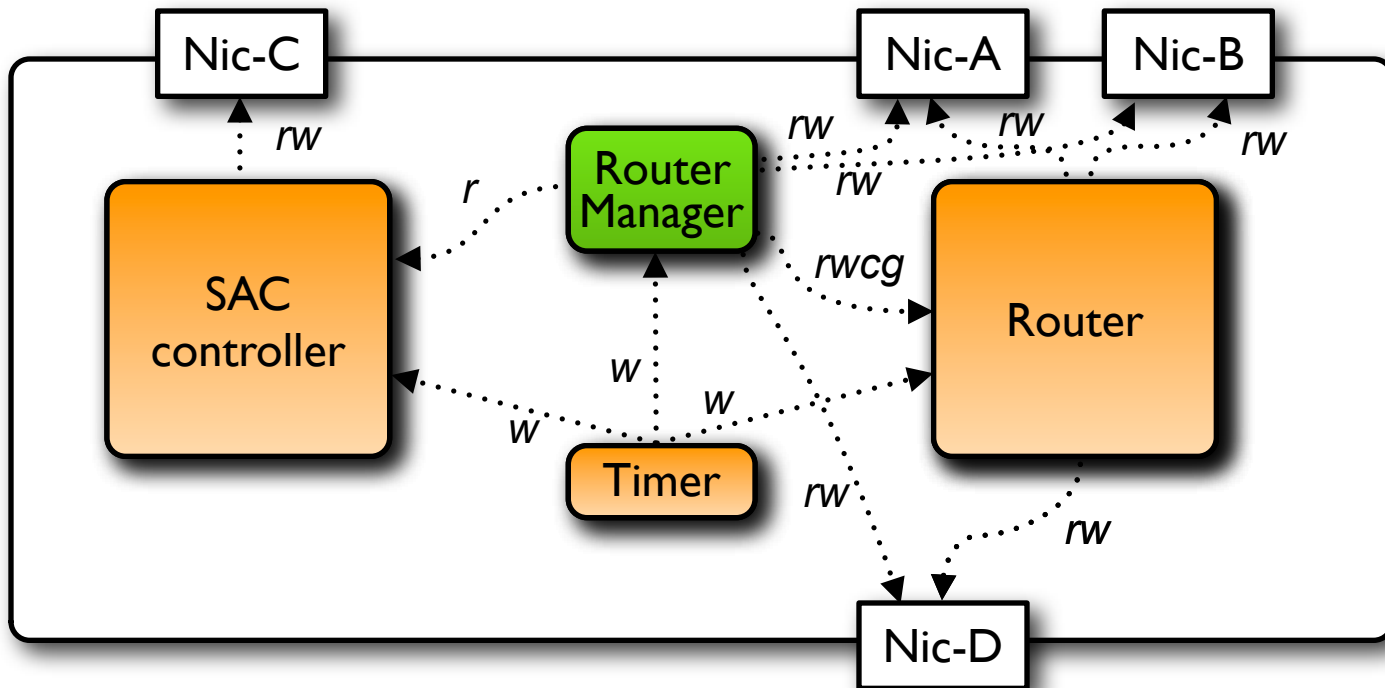
Now: how to set-up the system with this design?

Problem: reality is not that simple

Back to the example

This is what we agree on the whiteboard

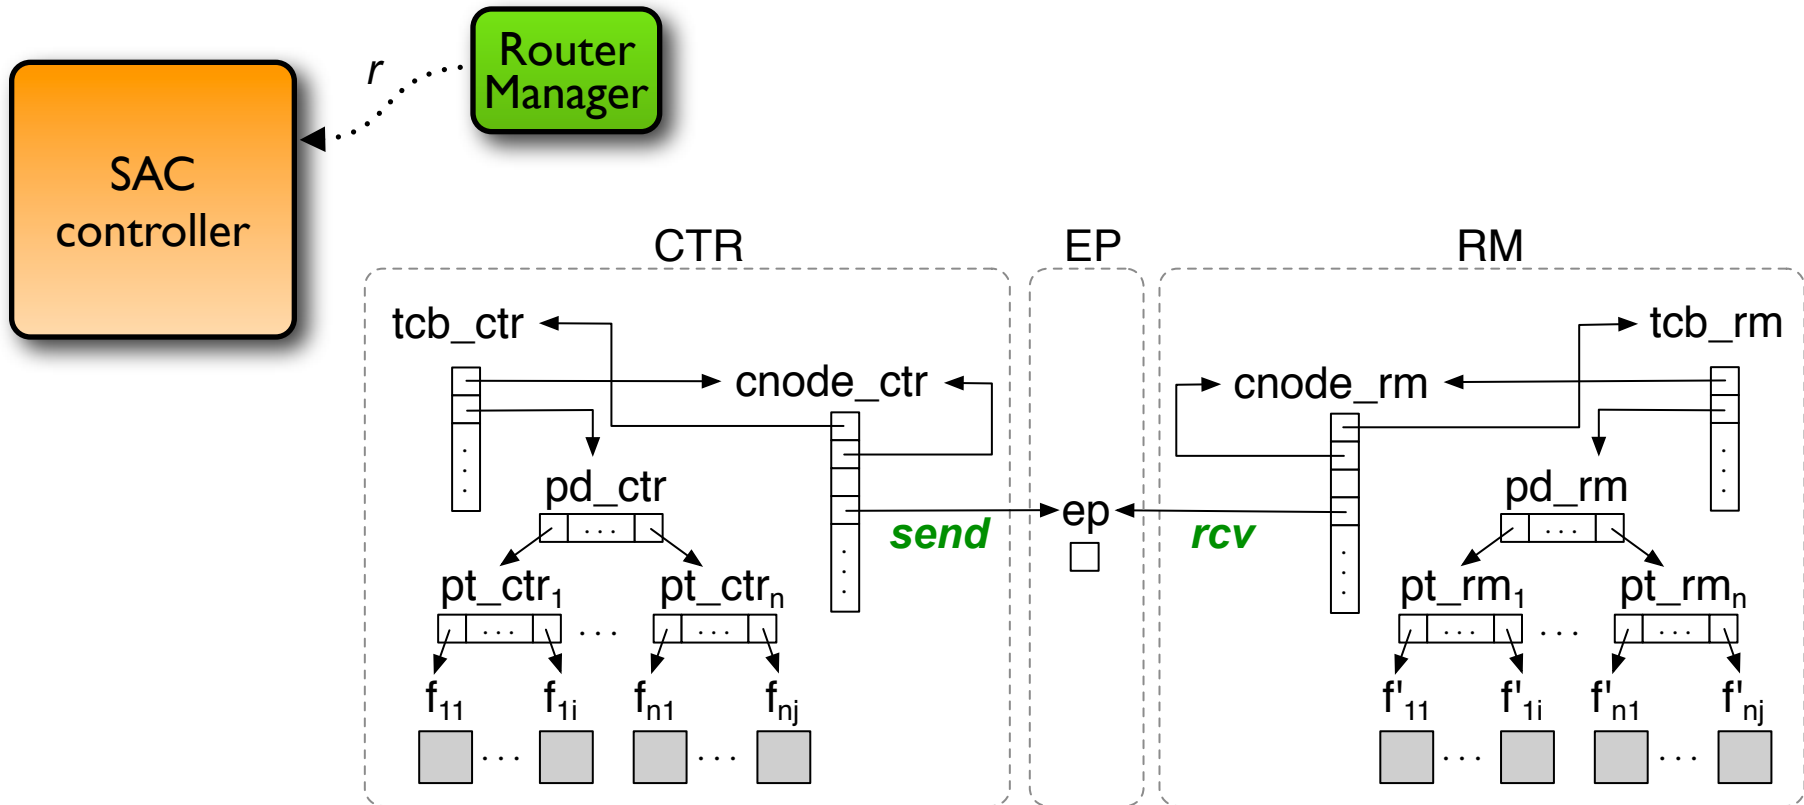
Now we need to implement this with actual kernel objects




Back to the example

This is what we agree on the whiteboard

Now we need to implement this with actual kernel objects

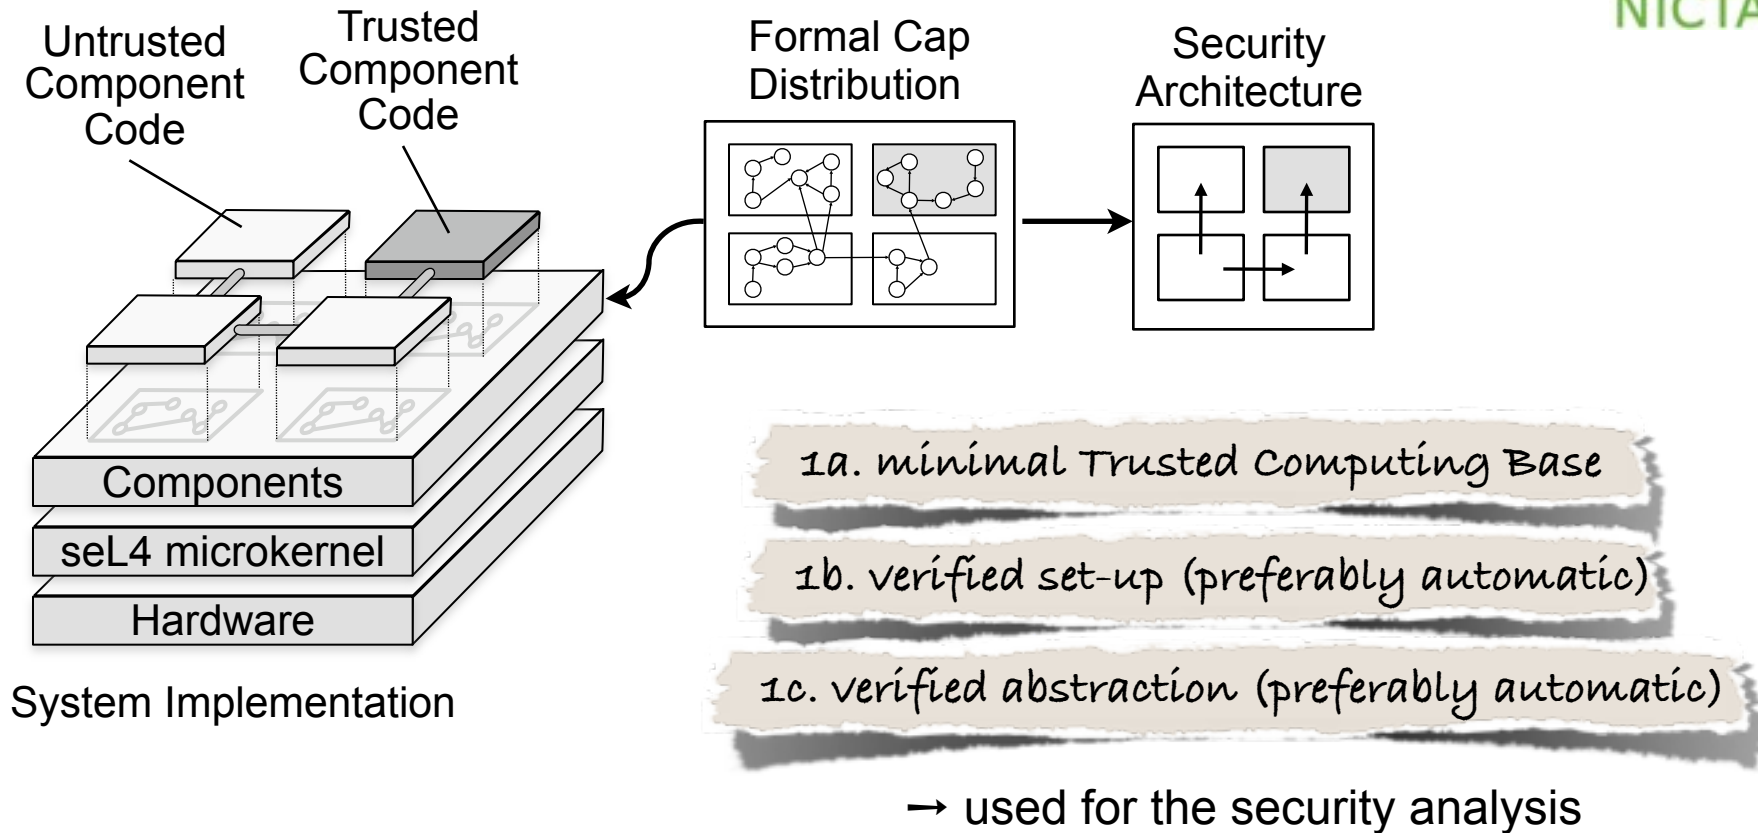


Every arrow is a capability!



capability distribution

Back to the general picture



capDL: capability distribution language

Example:

$obj1 \equiv Tcb[0 \mapsto CNodeCap\ 3, \dots]$

$obj3 \equiv CNode[302 \mapsto CNodeCap\ 9\ Read, \dots]$


Secure Systems "à la NICTA"

For 1 critical system:

- 1 desired security property P
- an interactive theorem prover
- a bit of patience

1. carefully design your system
2. prove that the design enforces P
3. prove correctness of the TCB
4. prove isolation

Security Proof




System Implementation

Theorem: $s_0 \xrightarrow{*} s \Rightarrow \mathcal{P}(s)$

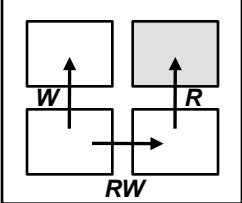
Lemma sacSecurity:

$(\text{SAC-startup} \xrightarrow{*} s) \Rightarrow$
 $\neg \text{is_contaminated } s \text{ NicA}$

Security Proof



Theorem: $s_0 \xrightarrow{*} s \Rightarrow \mathcal{P}(s)$


Where: $s_0 \equiv$ 

The diagram shows a 2x2 grid of boxes. The top-left box has an upward arrow labeled 'W'. The top-right box has an upward arrow labeled 'R'. The bottom-left box has a rightward arrow labeled 'RW'. The bottom-right box is shaded gray.

System Implementation


```
RM_id      -> Some ({cap_r_to_SAC_C, ...}, not_contaminated)
SAC_C_id   -> Some ({cap_rw_to_NIC_C, ...}, not_contaminated)
NIC_A_id   -> Some ( {}, not_contaminated)
NIC_B_id   -> Some ( {}, contaminated)
```

Security Proof



System Implementation

Theorem: $s_0 \xrightarrow{*} s \Rightarrow \mathcal{P}(s)$




$$s \rightarrow s' \equiv s \xrightarrow{t} s' \vee s \xrightarrow{u} s'$$

$s \xrightarrow{t} s' \equiv$ let $tc \in \text{trusted_component}(s)$ in
 let $\text{prg} = \text{program}(tc)$ in
 let $pc = \text{program_counter}(c, s)$ in
 let $i = \text{inst}(\text{prg}, pc)$ in
 $\text{step}(tc, s, i, s')$


```
RM_prg ≡
[ SysOp (SysRead cap_r_to_SAC_C),
  SysOp (SysRemoveAll cap_C_to_R),
  SysOp (SysDelete cap_C_to_R),
  ... ]
```

Security Proof



System Implementation

Theorem: $s_0 \xrightarrow{*} s \Rightarrow \mathcal{P}(s)$

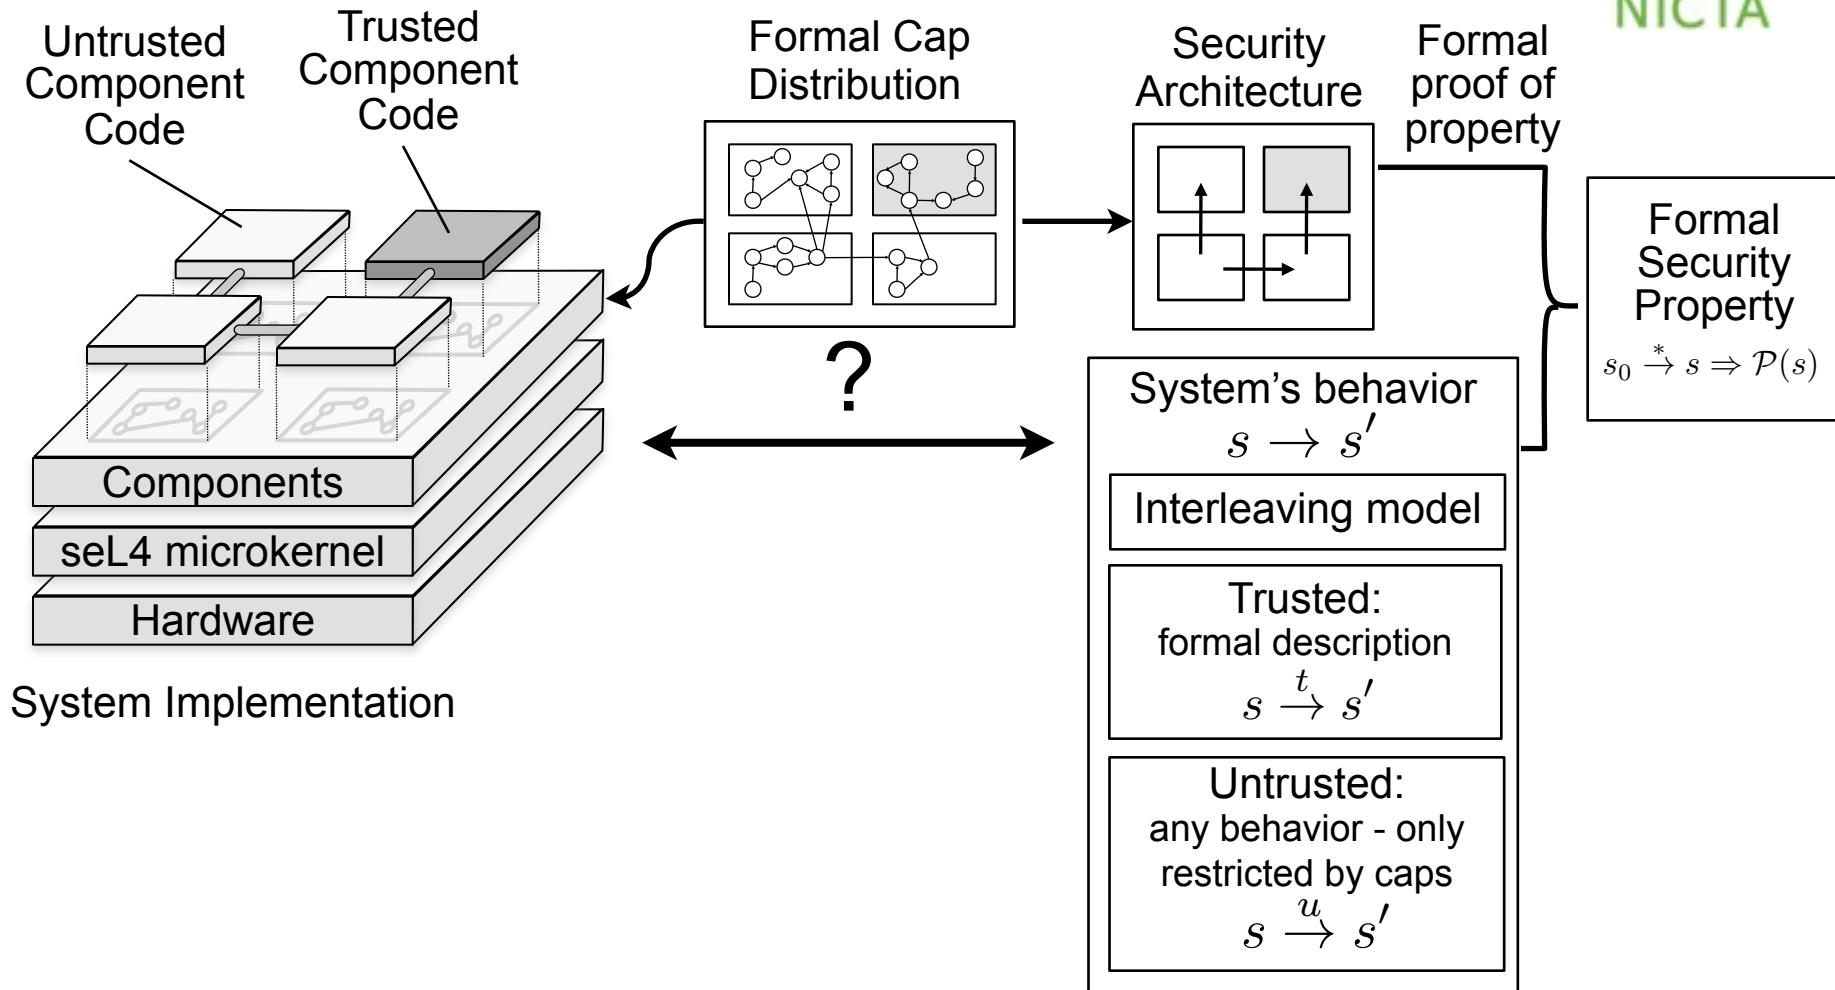


$$s \rightarrow s' \equiv s \xrightarrow{t} s' \vee s \xrightarrow{u} s'$$

$$s \xrightarrow{t} s' \equiv \text{let } tc \in \text{trusted_component}(s) \text{ in} \\ \text{let } prg = \text{program}(tc) \text{ in} \\ \text{let } pc = \text{program_counter}(c, s) \text{ in} \\ \text{let } i = \text{inst}(prg, pc) \text{ in} \\ \text{step}(tc, s, i, s')$$

$$s \xrightarrow{u} s' \equiv \text{let } uc \in \text{untrusted_components}(s) \text{ in} \\ \text{step}(uc, s, \text{any_inst}, s')$$

Security Proof



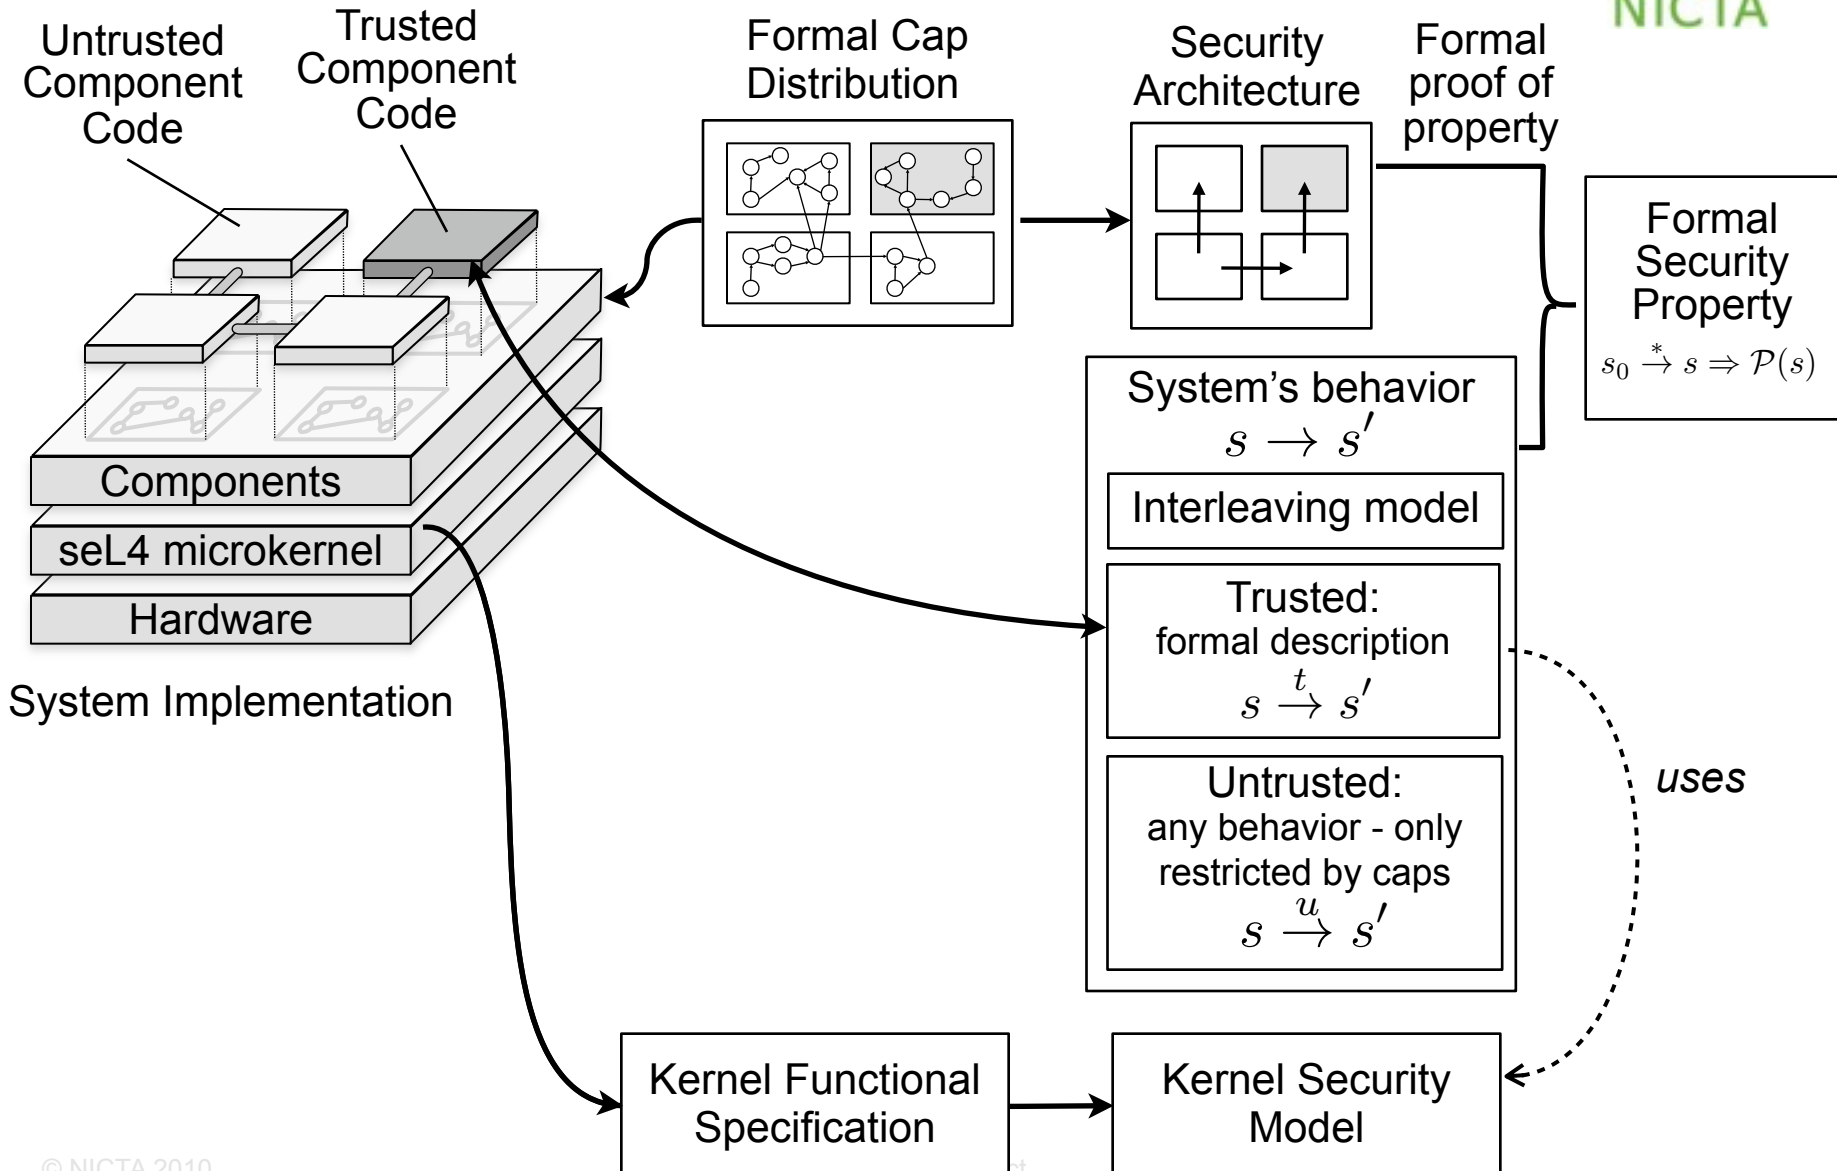
Secure Systems "à la NICTA"

For 1 critical system:

- 1 desired security property P
- an interactive theorem prover
- a bit of patience

1. carefully design your system
2. prove that the design enforces P
3. prove correctness of the TCB
4. prove isolation

Verified TCB



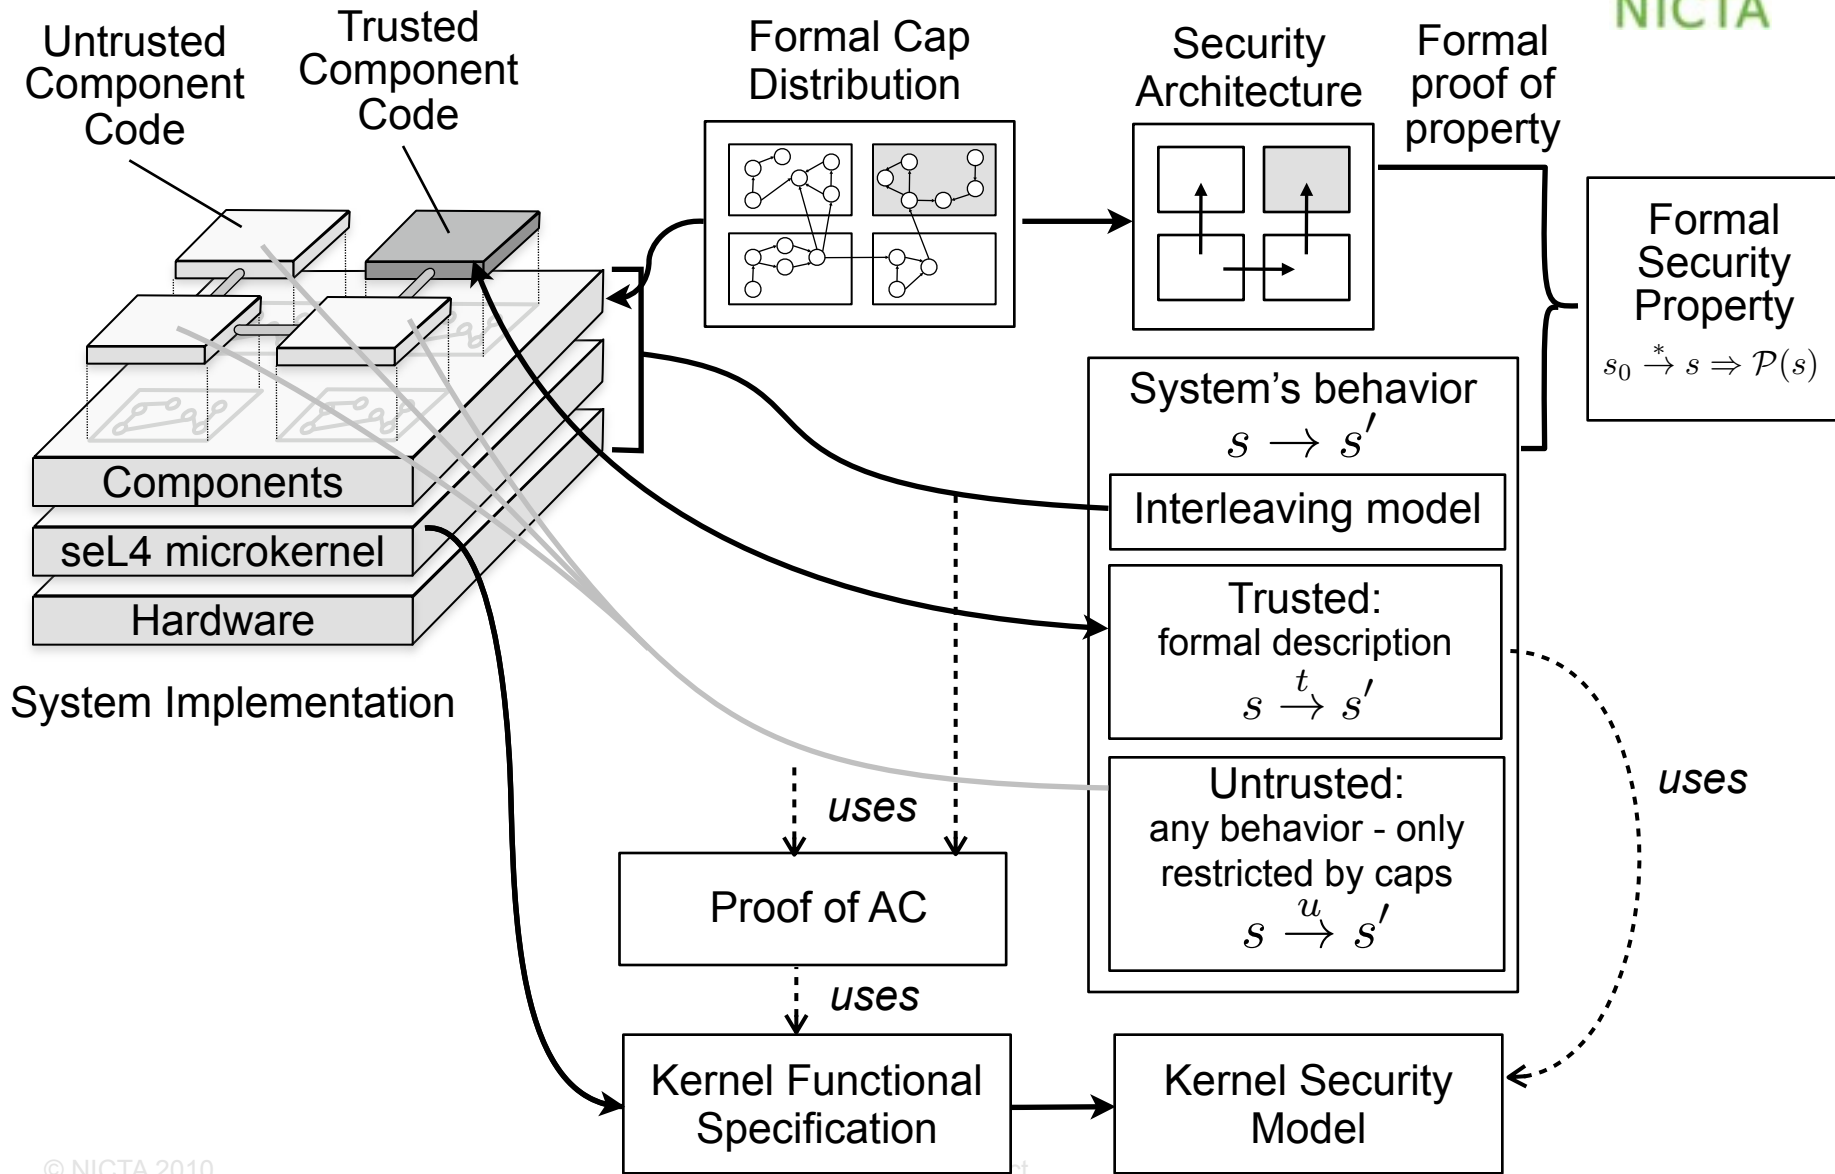
Secure Systems "à la NICTA"

For 1 critical system:

- 1 desired security property P
- an interactive theorem prover
- a bit of patience

1. carefully design your system
2. prove that the design enforces P
3. prove correctness of the TCB
4. prove isolation

Proof of access control



What is AC good for?



What is AC good for?





Examples

- R does not write to NicB if it does not have a write capability to it
- R does not change RM's program counter

Question: for all operation op s.t.

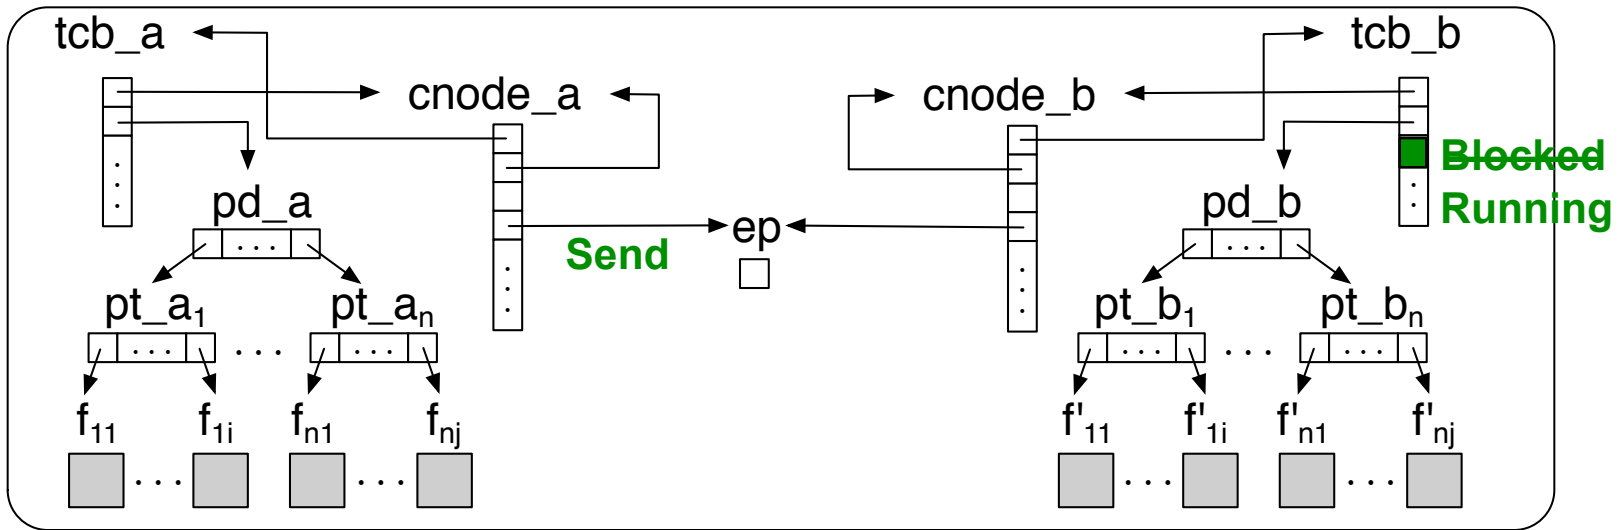
$$s \xrightarrow{op} s'$$

what is allowed to change in s' ?



Example


If *op* is *set_thread_state tcb_b v*
 If *tcb_a* is running in state *s* where *s* is:



Then in which condition may *tcb_b* change
 and what is allowed to change?

Example

If *op* is *set_thread_state tcb_b v*
 If *tcb_a* is running in state *s* where *s* is:




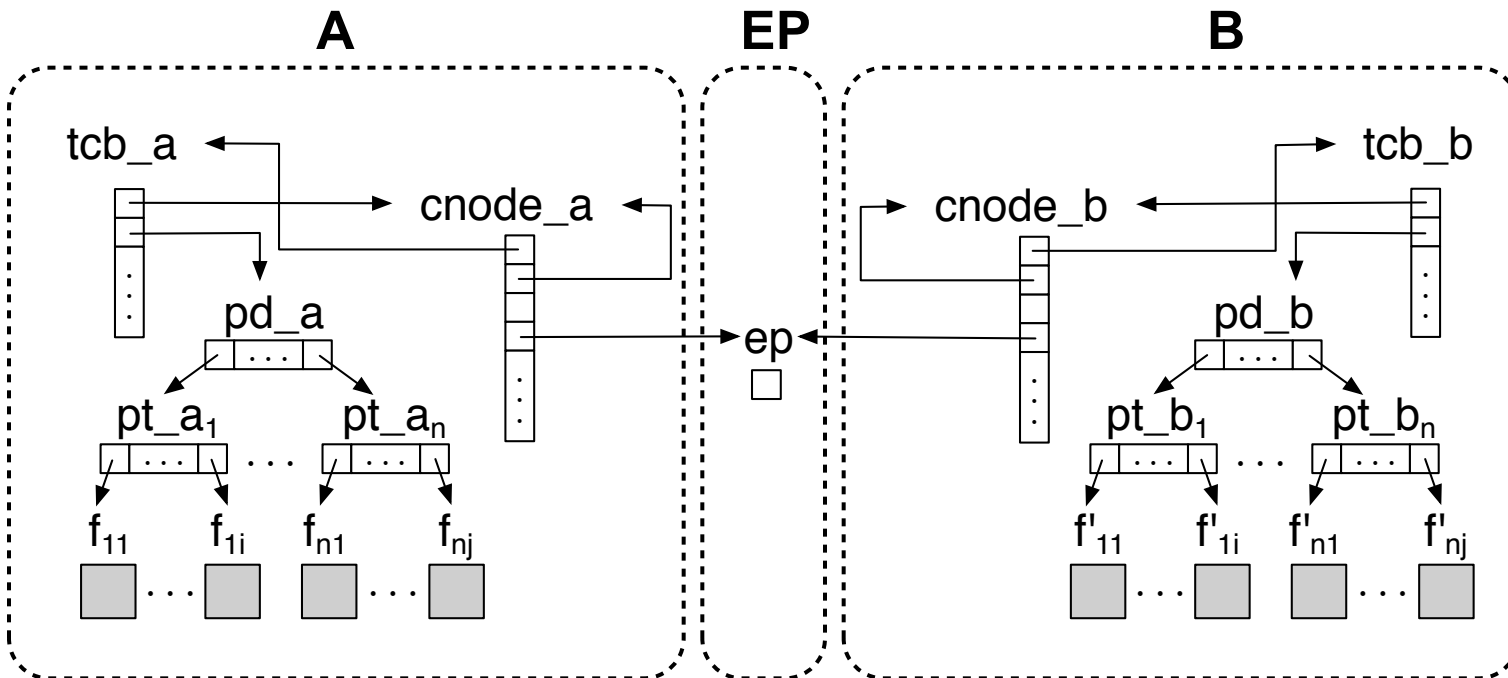
Obvious (but wrong) solution:

only the thread-state field of *tcb_b* is allowed to change
 and only under the following conditions:

- *tcb_a* has a capability to change the state of *tcb_b*
- or *tcb_a* is the parent of *tcb_b*
- or *tcb_a* is the parent of *tcb_b* and *tcb_b* is waiting on in state *s*
- or *tcb_a* is the parent of *tcb_b* and *tcb_b* is waiting on in state *s* and the untyped region containing *tcb_b*, in state *s*
- or ...

Policy closely depends on state

Solution: Labelling



Given by user

State must be subset a of G

G is subjective: current label contains the (untrusted) running thread

We prove:

1. **Graph preservation (authority confinement)**
2. **Access control at the label level**

Solution: Labelling

If **A** is the running label in **G**

then for **any operation *op*** that changes **s** to **s'**,

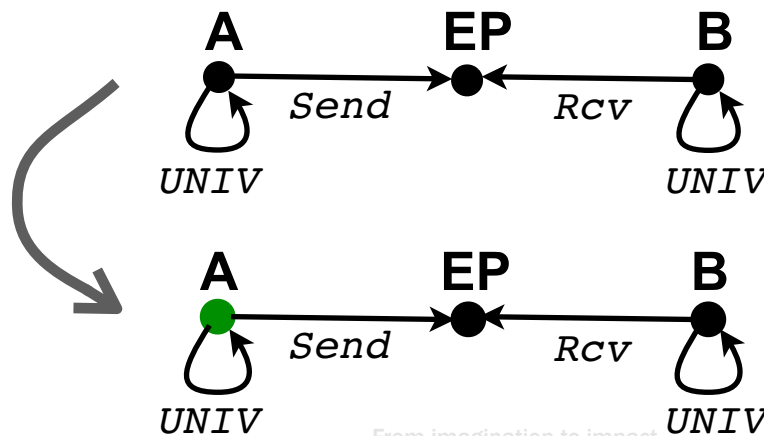
for any object **obj** of label **B**,

obj can only be changed if **A=B** or in 4 small precise cases, as:

“**obj** is a TCB blocked on an endpoint of label **EP**,

and $(A, \text{Send}, EP) \subseteq G$

and only the thread-state of **obj** can be changed, to Running”



We prove:

1. **Graph preservation (authority confinement)**
2. **Access control at the label level**

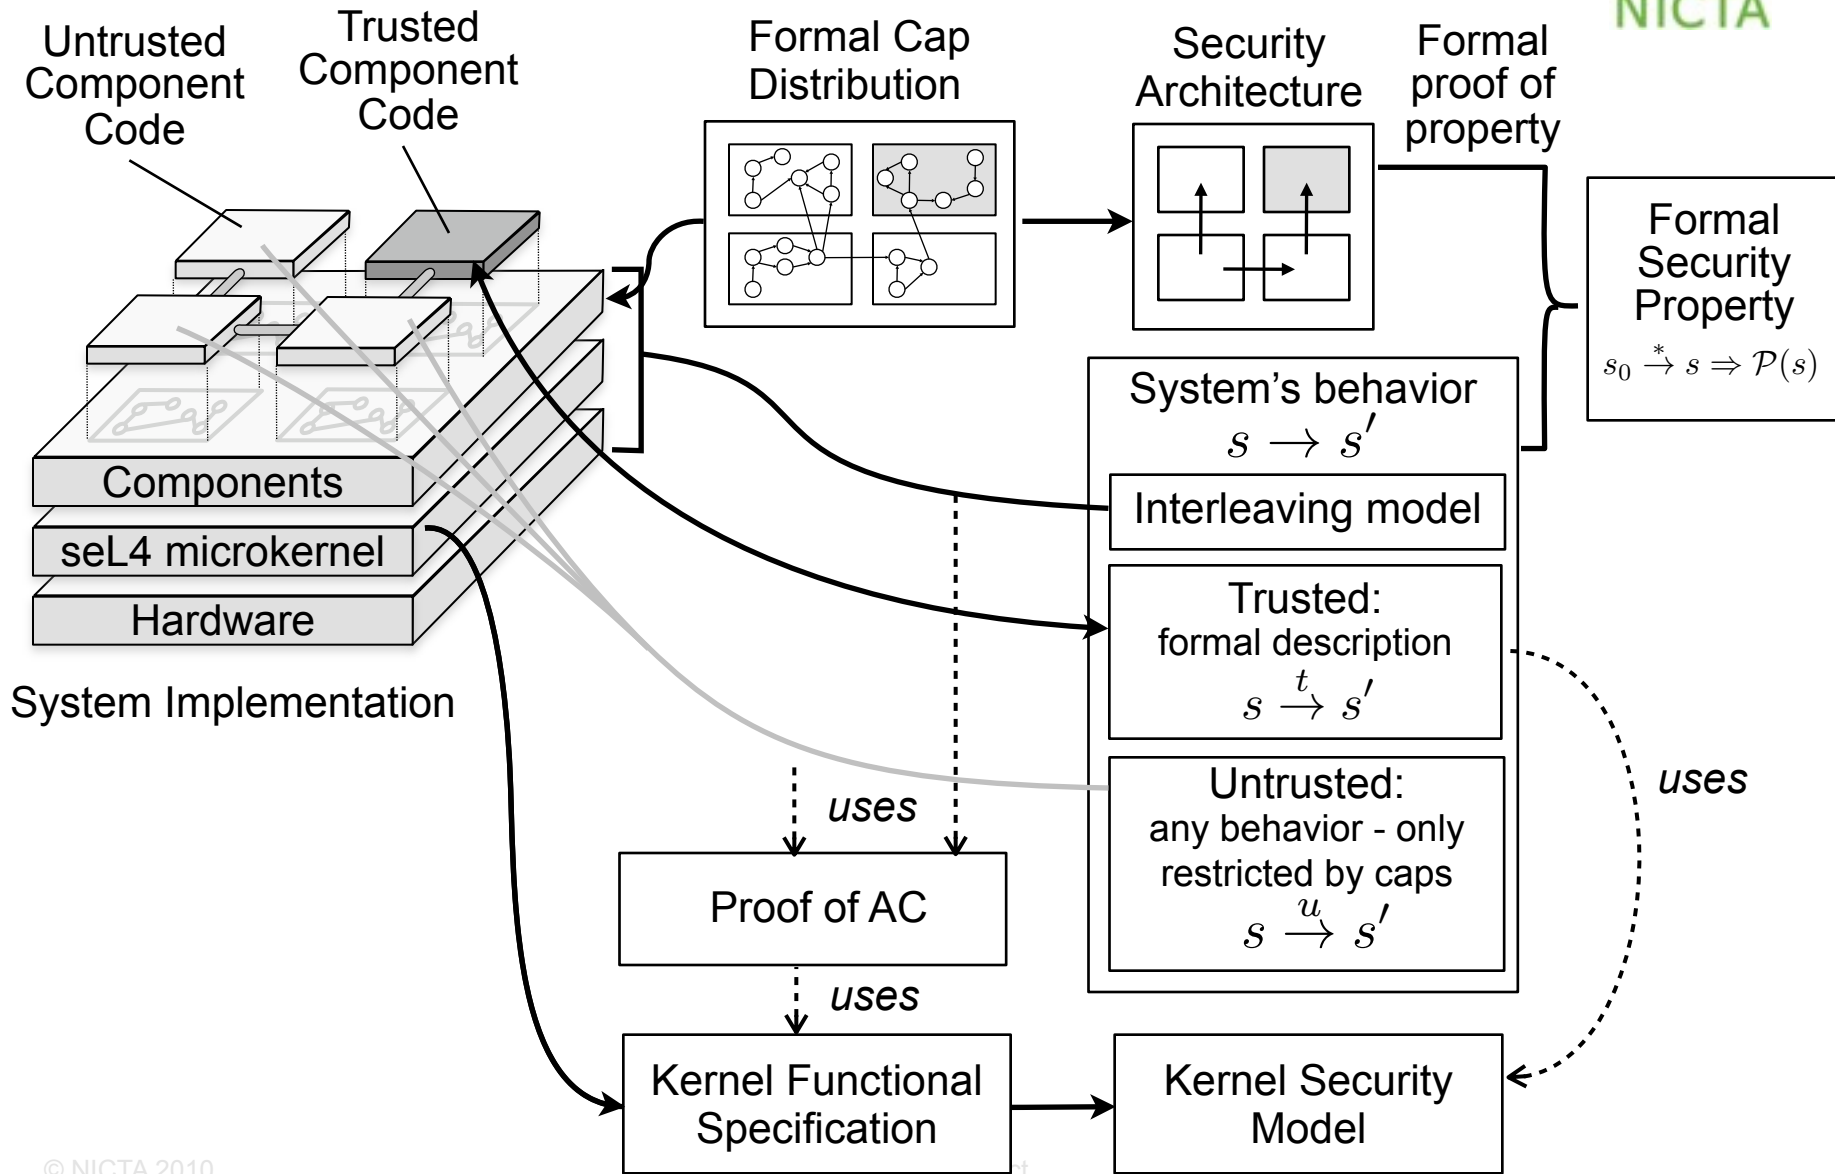
Secure Systems "à la NICTA"

For 1 critical system:

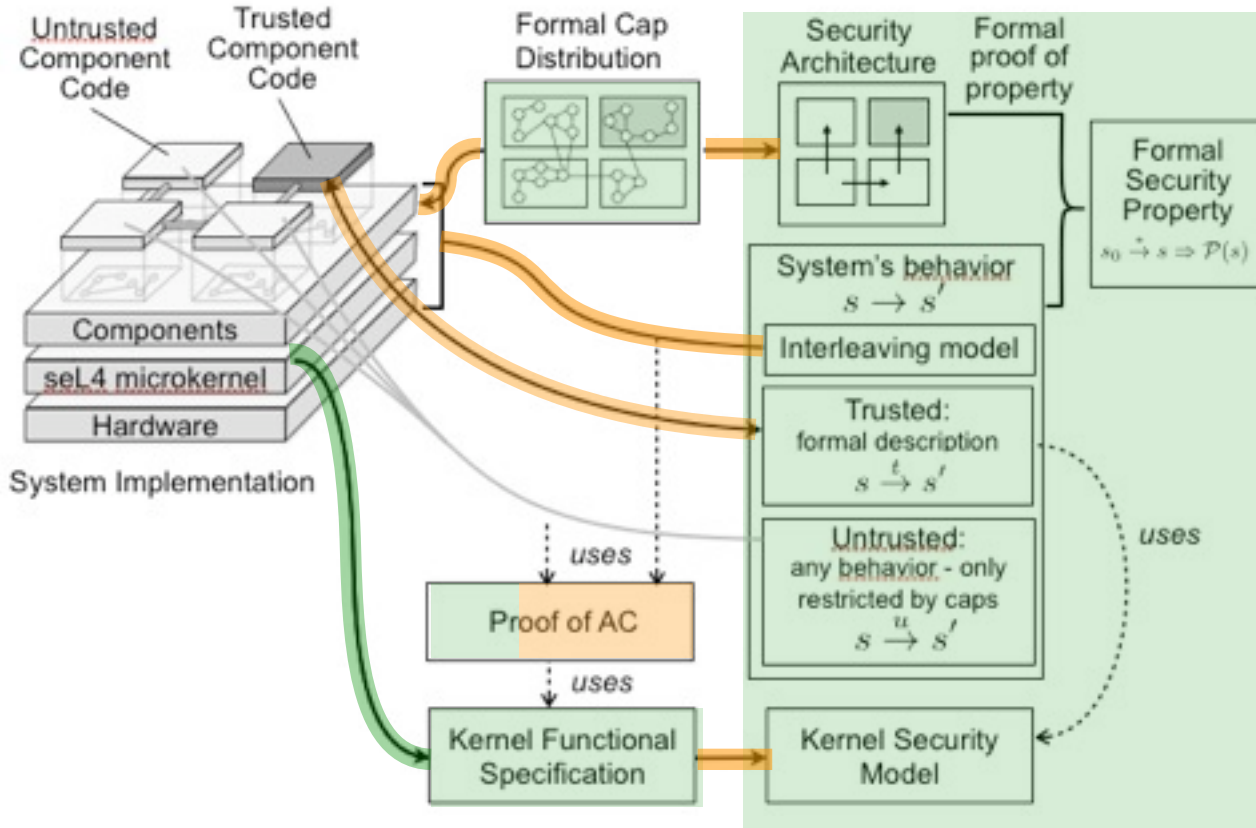
- 1 desired security property P
- an interactive theorem prover
- a bit of patience

1. carefully design your system
2. prove that the design enforces P
3. prove correctness of the TCB
4. prove isolation

Conclusion



Conclusion



Results so far:

- Case study sec. proof
- seL4 enforces integrity
- capDL
- seL4 correctness proof
- certification: ST
- seL4 binary and formal spec released

Challenges:

- | | | | |
|---|---|---|--|
| <p>Automation
verified code generator
from high level code</p> | <p>Integration
system trace reasoning,
concurrency</p> | <p>Confidentiality
preserved by
refinement</p> | <p>Proof engineering</p> <ul style="list-style-type: none"> • refactoring • efficient proof rerun |
|---|---|---|--|



Questions?