

# A Verified OS Kernel. Now What?

Gerwin Klein



**Australian Government**  
Department of Communications,  
Information Technology and the Arts  
Australian Research Council

NICTA Members



Department of State and  
Regional Development



The University of Sydney



Queensland University of Technology



NICTA Partners

# The Team


---





# The Team

---



**L4** **Verified**

# L4 Verified

1 microkernel

8,700 lines of C

0 bugs\*

qed

\*conditions apply

## Windows

An exception 06 has occurred at 0028:C11B3ADC in VxD DiskTSD(03) + 00001660. This was called from 0028:C11B40C8 in VxD voltrack(04) + 00000000. It may be possible to continue normally.

- \* Press any key to attempt to continue.
- \* Press CTRL+ALT+RESET to restart your computer. You will lose any unsaved information in all applications.

Press any key to continue






# Windows Vista

## Stunning. Breakthrough. Entertaining.

HP TouchSmart PC and Microsoft Windows Vista deliver you a PC experience designed to fit wherever life happens.

Innovative solutions brought to you by:

Microsoft









# The Problem

---





# Small Kernels

## Small trustworthy foundation

- hypervisor, microkernel, nano-kernel, virtual machine, separation kernel, exokernel ...
- High assurance components in presence of other components

### seL4 API:

- IPC
- Threads
- VM
- IRQ
- Capabilities

## Untrusted

Legacy Apps

Linux Server

## Trusted

Sensitive App

Trusted Service

Hardware

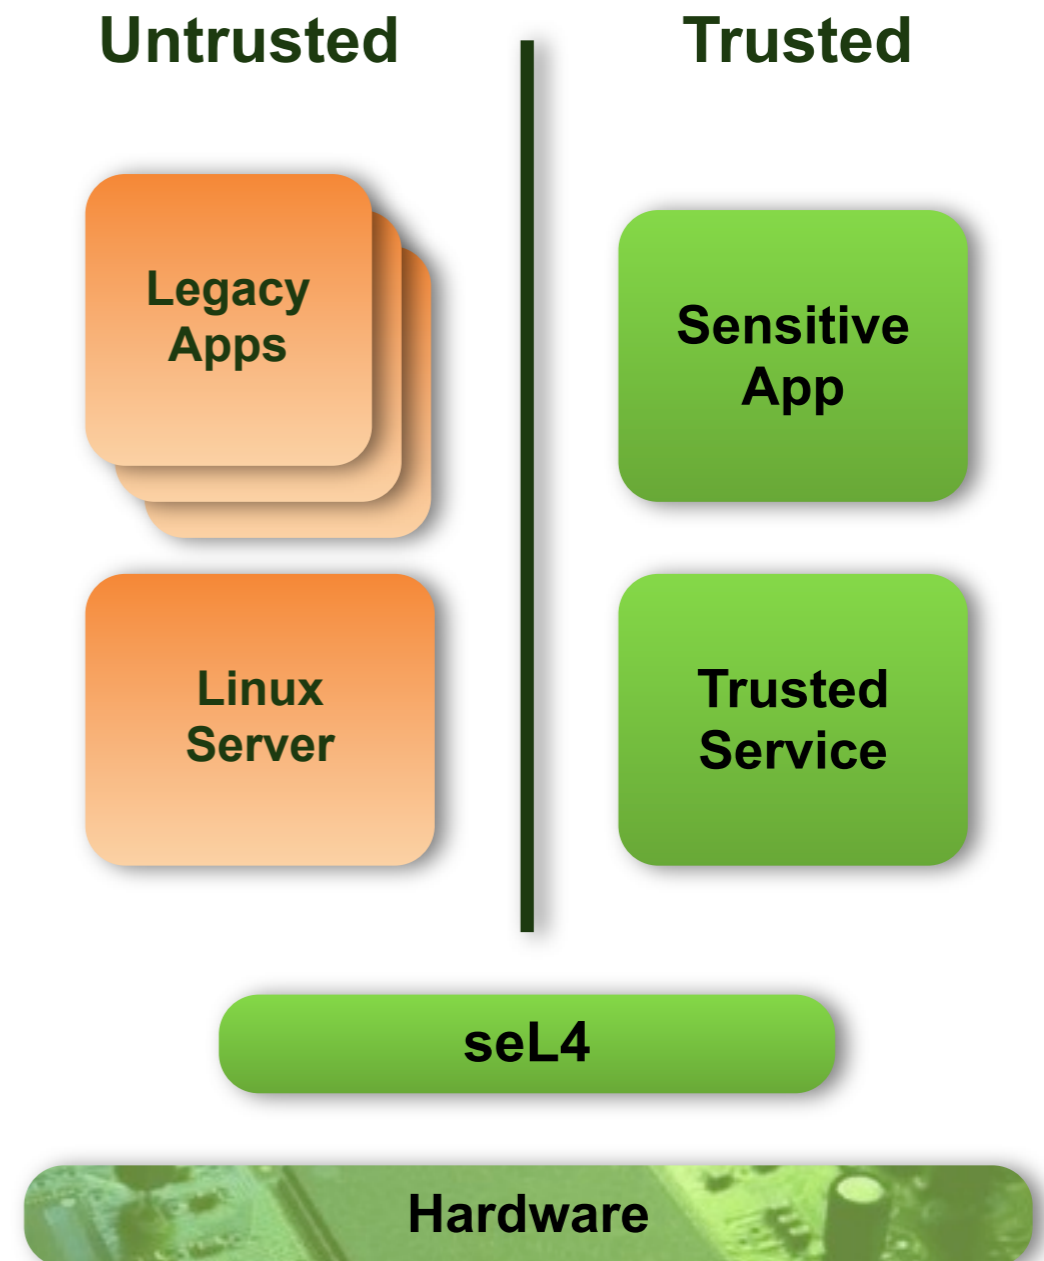
# Small Kernels

## Small trustworthy foundation

- hypervisor, microkernel, nano-kernel, virtual machine, separation kernel, exokernel ...
- High assurance components in presence of other components

### seL4 API:


- IPC
- Threads
- VM
- IRQ
- Capabilities



# The Proof




# The Proof



# Functional Correctness

---



# Functional Correctness

**What**

**definition**

```
schedule :: unit s_monad where  
schedule ≡ do  
  threads ← allActiveTCBs;  
  thread ← select threads;  
  switch_to_thread thread  
od  
OR switch_to_idle_thread
```

**Specification**

**Proof**



**Code**



# Functional Correctness



What

Specification

definition

```
schedule :: unit s_monad where
schedule ≡ do
  threads ← allActiveTCBs;
  thread ← select threads;
  switch_to_thread thread
od
OR switch_to_idle_thread
```

Proof

How

```
void
schedule(void) {
  switch ((word_t)ksSchedulerAction) {
    case (word_t)SchedulerAction_ResumeCurrentThread:
      break;

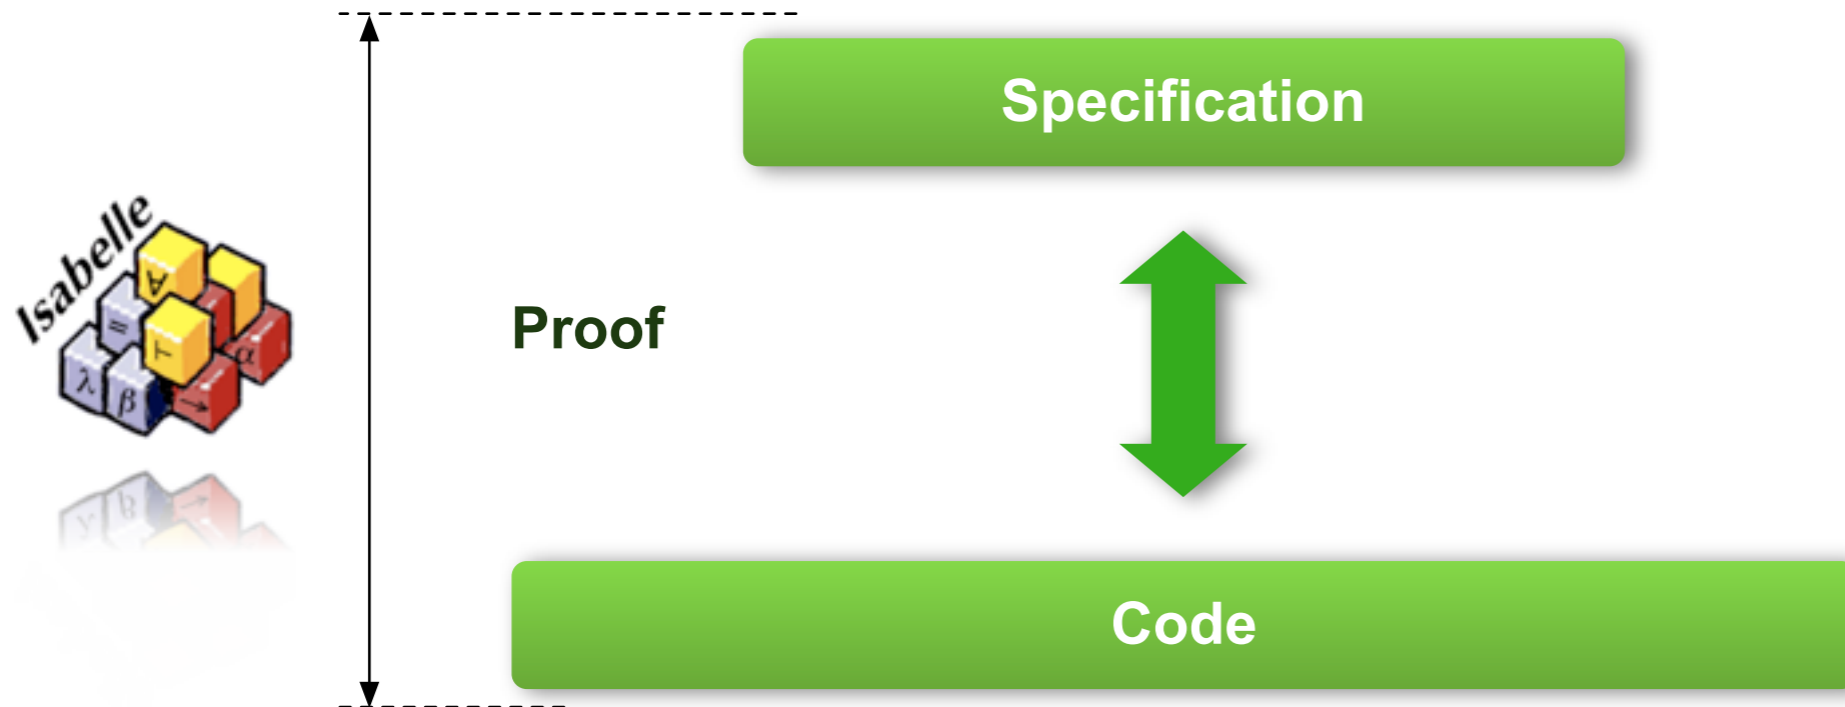
    case (word_t)SchedulerAction_ChooseNewThread:
      chooseThread();
      ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
      break;

    default: /* SwitchToThread */
      switchToThread(ksSchedulerAction);
      ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
      break;
  }
}

void
chooseThread(void) {
  prio_t prio;
  tcb_t *thread, *next;
```



# \*conditions apply



\*conditions apply



Expectation

Specification

Proof




Code

Assumptions



\*conditions apply



**Assume correct:**

- compiler + linker (wrt. C op-sem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

**Assumptions**



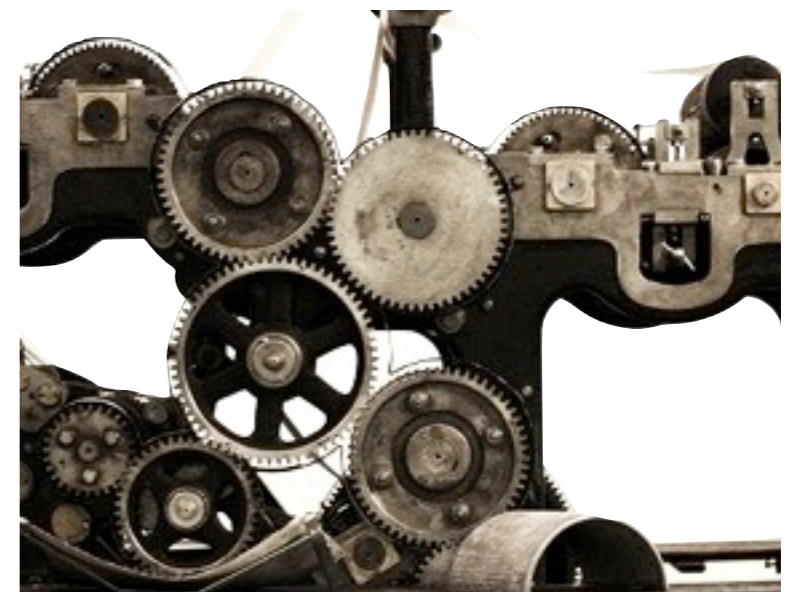
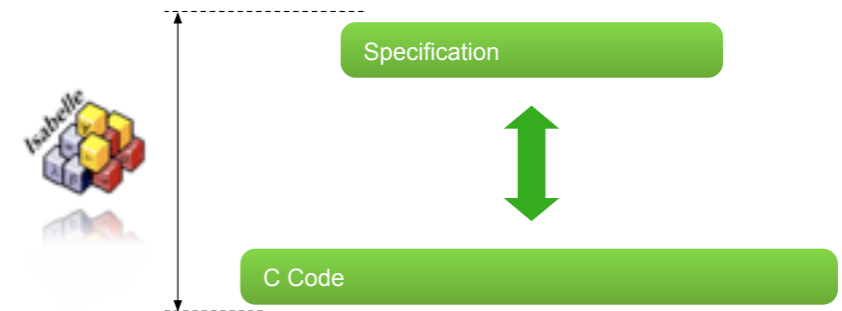
# Implications

## Execution always defined:

- no null pointer de-reference
- no buffer overflows
- no code injection
- no memory leaks/out of kernel memory
- no div by zero, no undefined shift
- no undefined execution
- no infinite loops/recursion

## Not implied:

- “secure” (define secure)
- zero bugs from expectation to physical world
- covert channel analysis





# Implications



## Execution always defined:

- no null pointer de-reference
- no buffer overflows
- no code injection
- no memory leaks/out of kernel memory
- no div by zero, no undefined shifts
- no undefined execution
- no infinite loops/recursion

## Not implied:

- “secure” (define secure)
- zero bugs from expectation to physical world
- covert channel analysis

**THE H SECURITY** The H open source security In association with

Last 7 days News Archive Features Forums Newsletter RSS

14 August 2009, 12:14 << previous | next >>

### Critical vulnerability in the Linux kernel affects all versions since 2001

Google security specialists Tavis Ormandy and Julien Tiennes report that a critical security vulnerability in the [Linux kernel](#) affects all versions of 2.4 and 2.6 since 2001, on all architectures. The vulnerability enables users with limited rights to get root rights on the system. The cause is a NULL pointer dereference in connection with the initialisation of sockets for rarely used protocols.



# Implications

## Execution always defined:

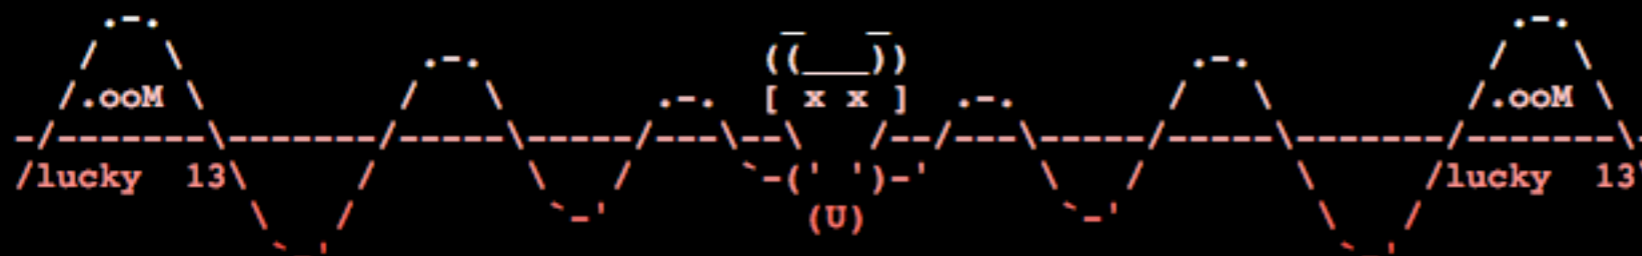
- no null pointer de-reference
- no buffer overflows



Specification



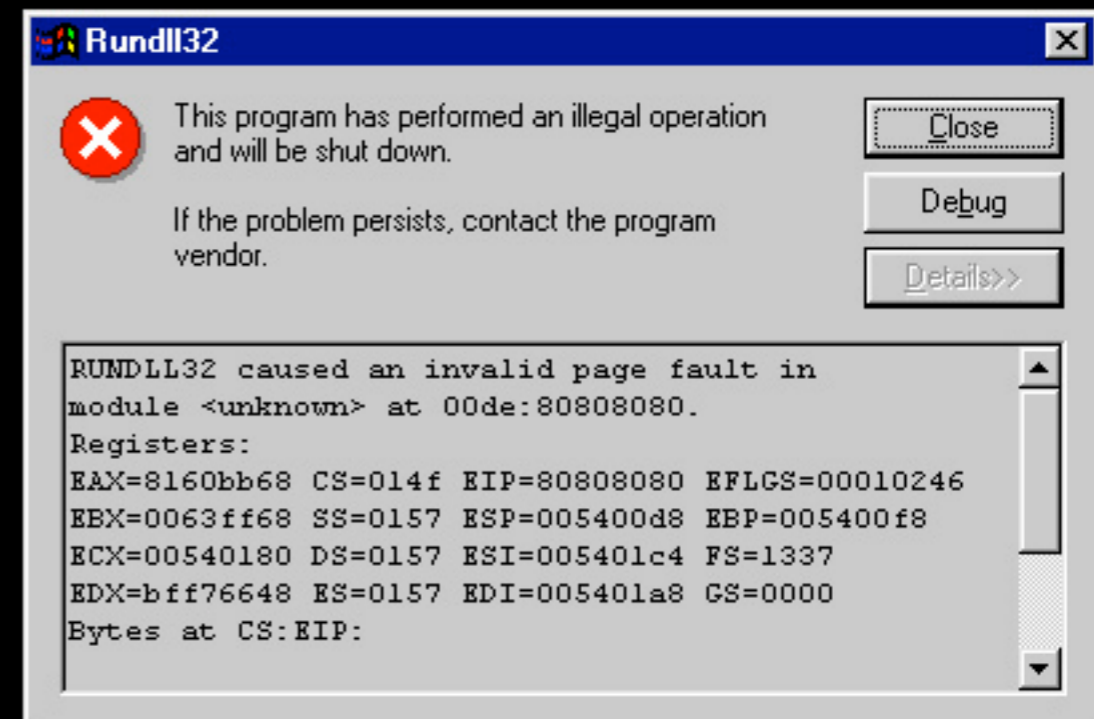
C Code



## The Tao of Windows Buffer Overflow

*as taught by*  
**DiIDog**  
cDc Ninja Strike Force  
9-dan of the Architecture  
Sensei of the Undocumented Opcode

[Begin](#)




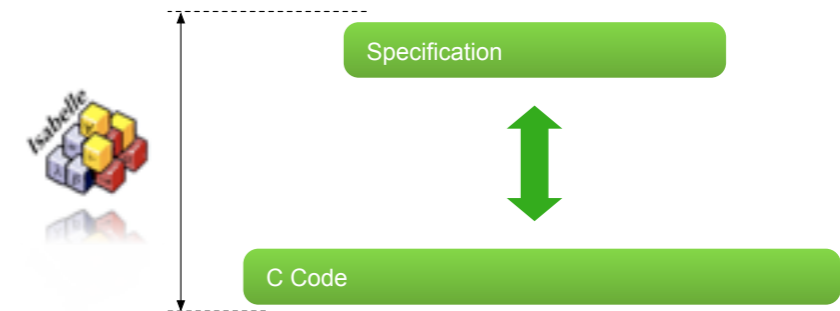
# Implications

## Execution always defined:

- no null pointer de-reference
- no buffer overflows
- no code injection
- no memory leaks/out of kernel memory
- no div by zero, no undefined shift
- no undefined execution
- no infinite loops/recursion


## Not implied:

- “secure” (define secure)
- zero bugs from expectation to physical world
- covert channel analysis





# Proof Architecture



# Proof Architecture

---

Specification

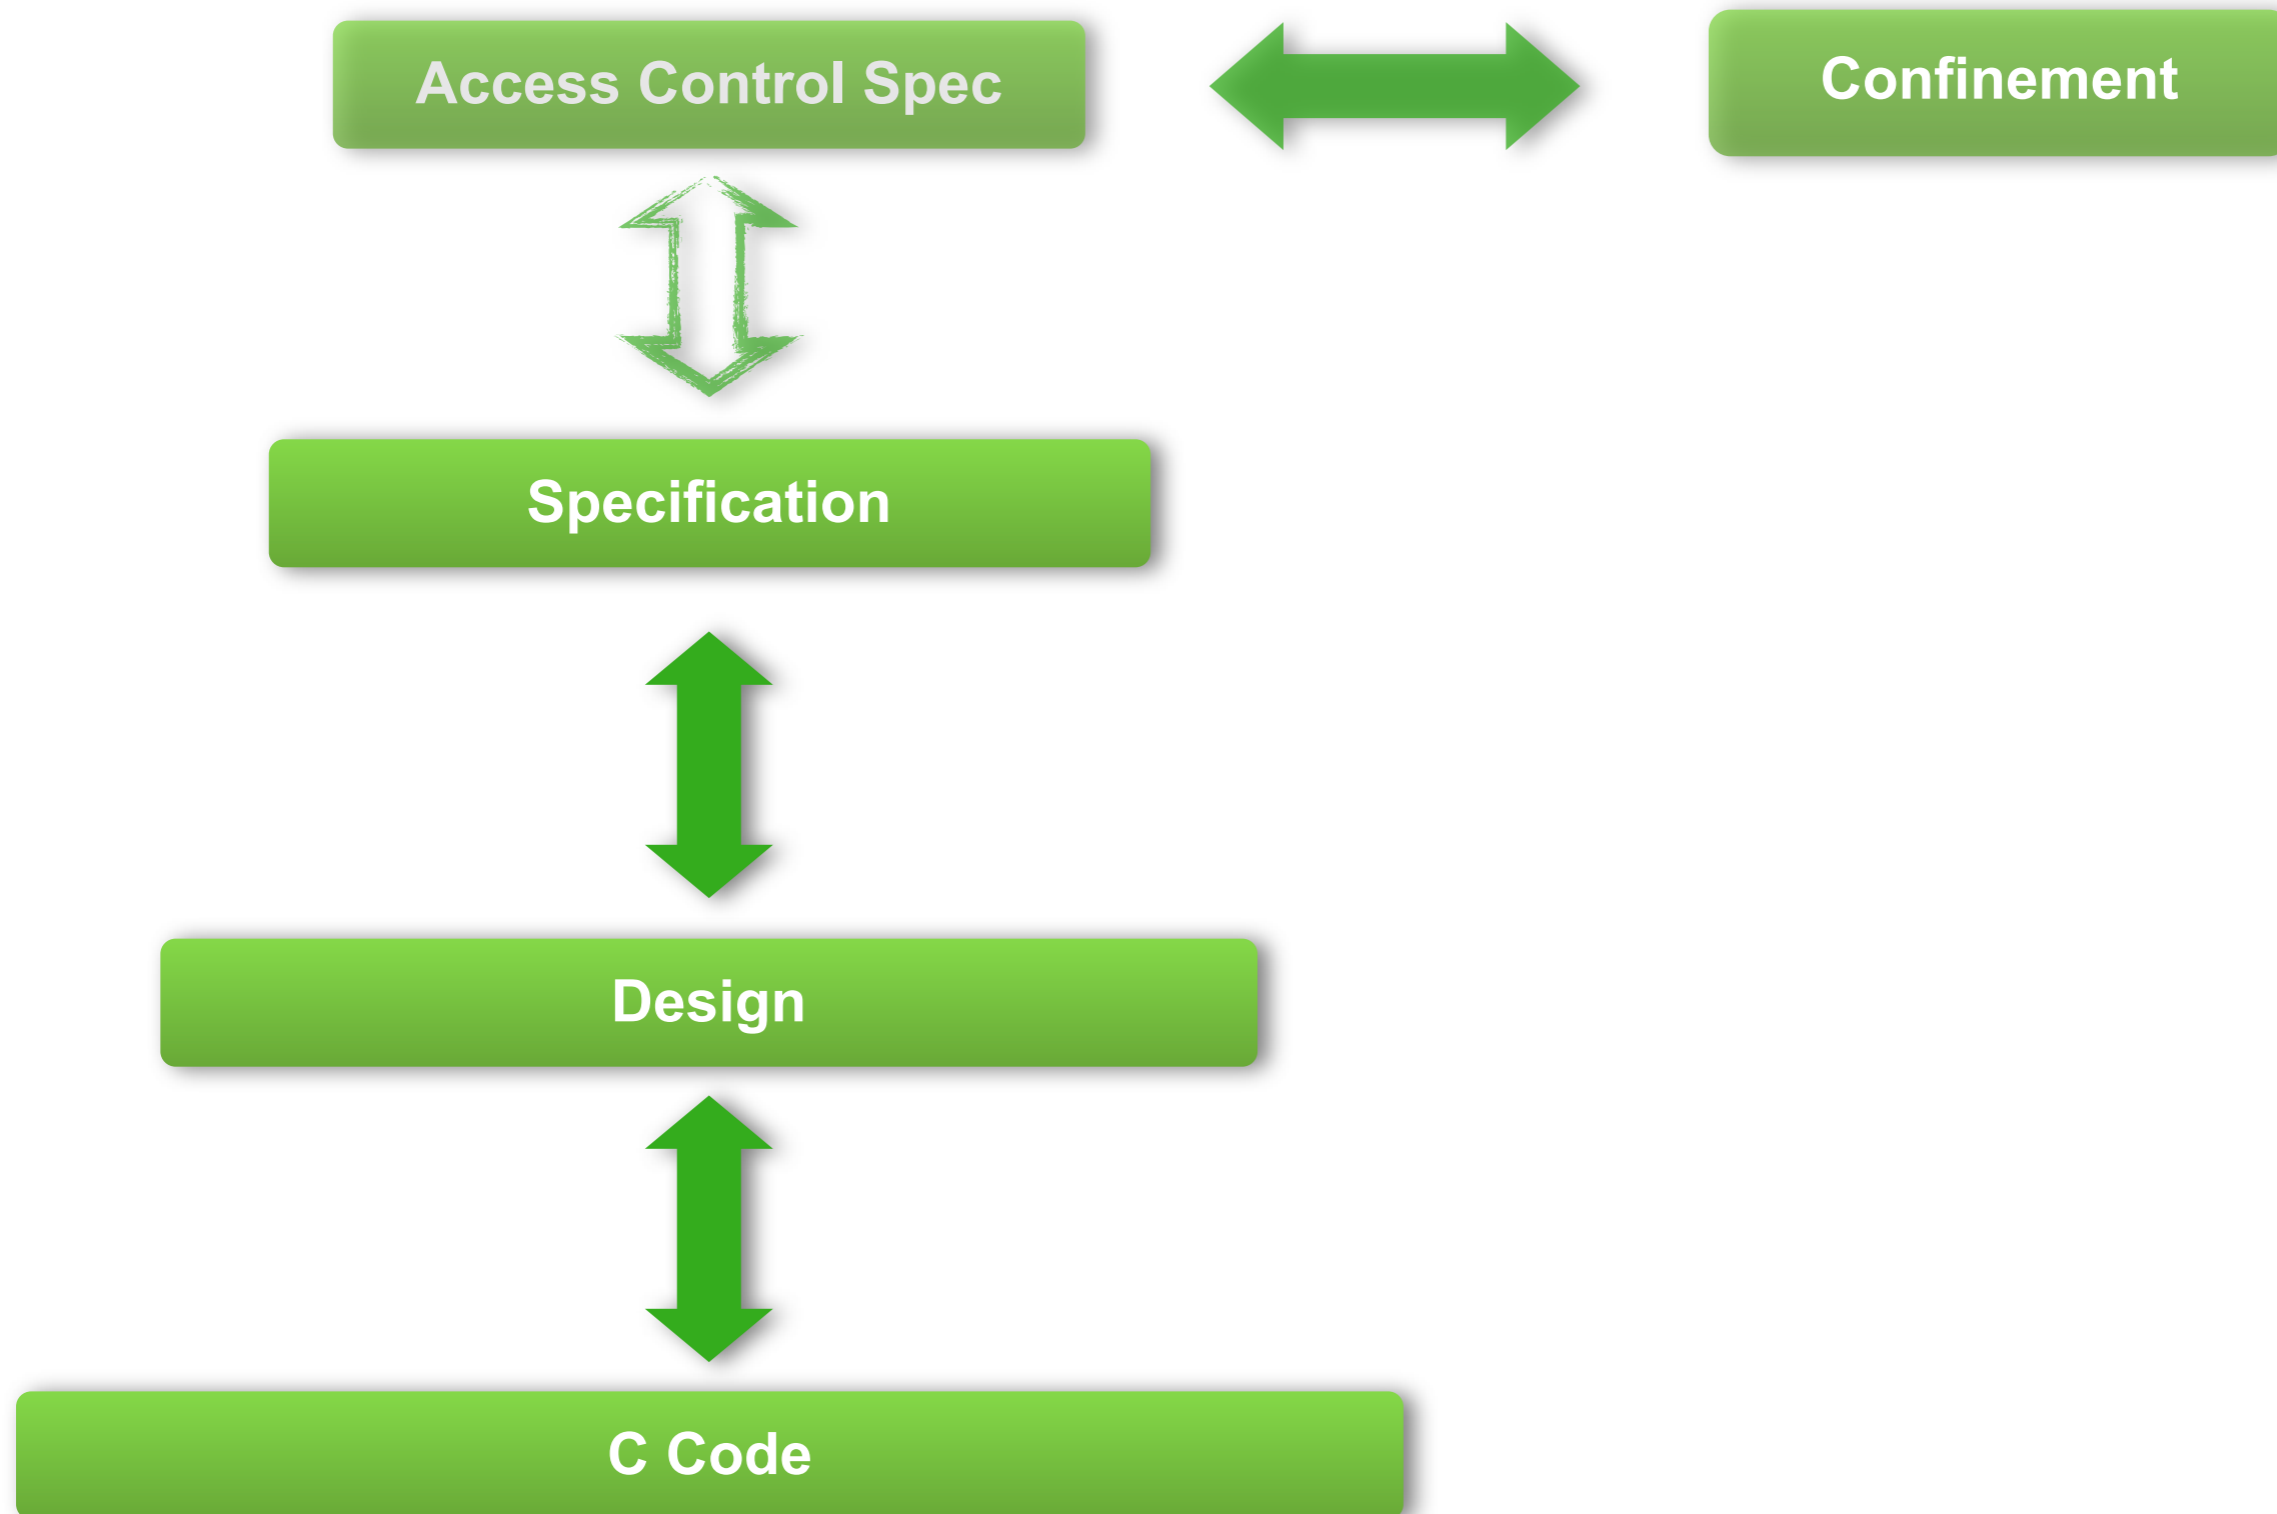


Design

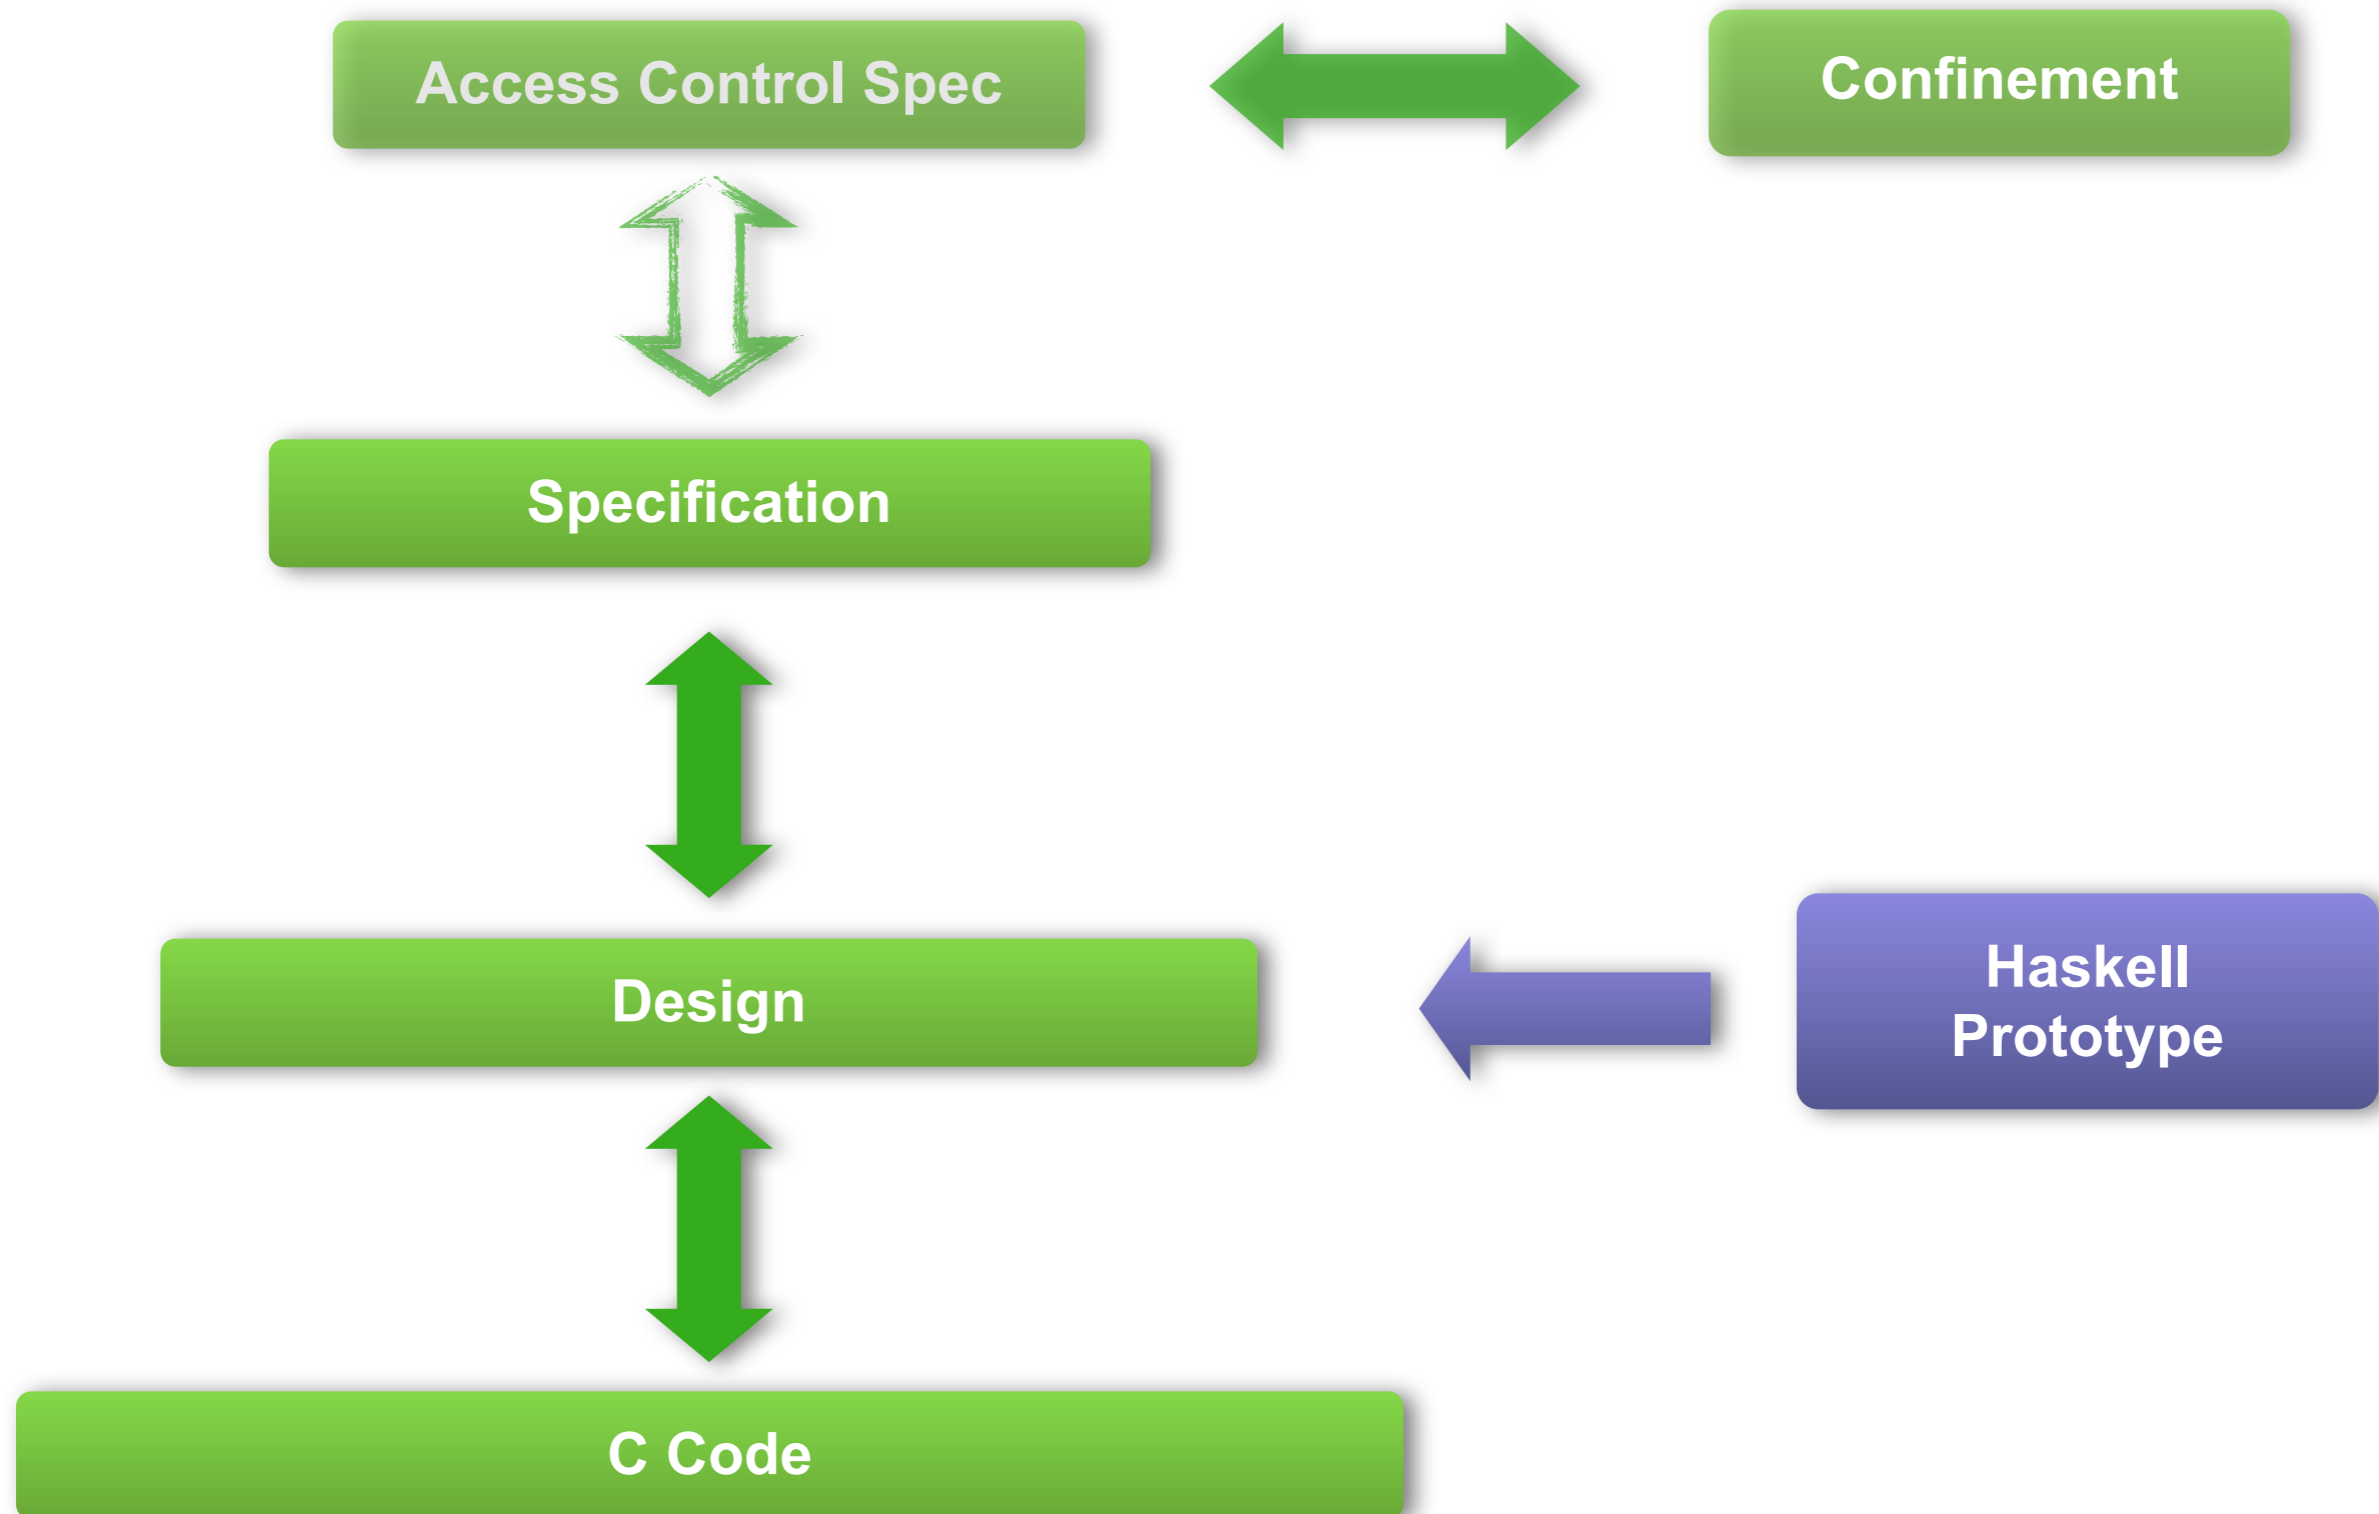


C Code

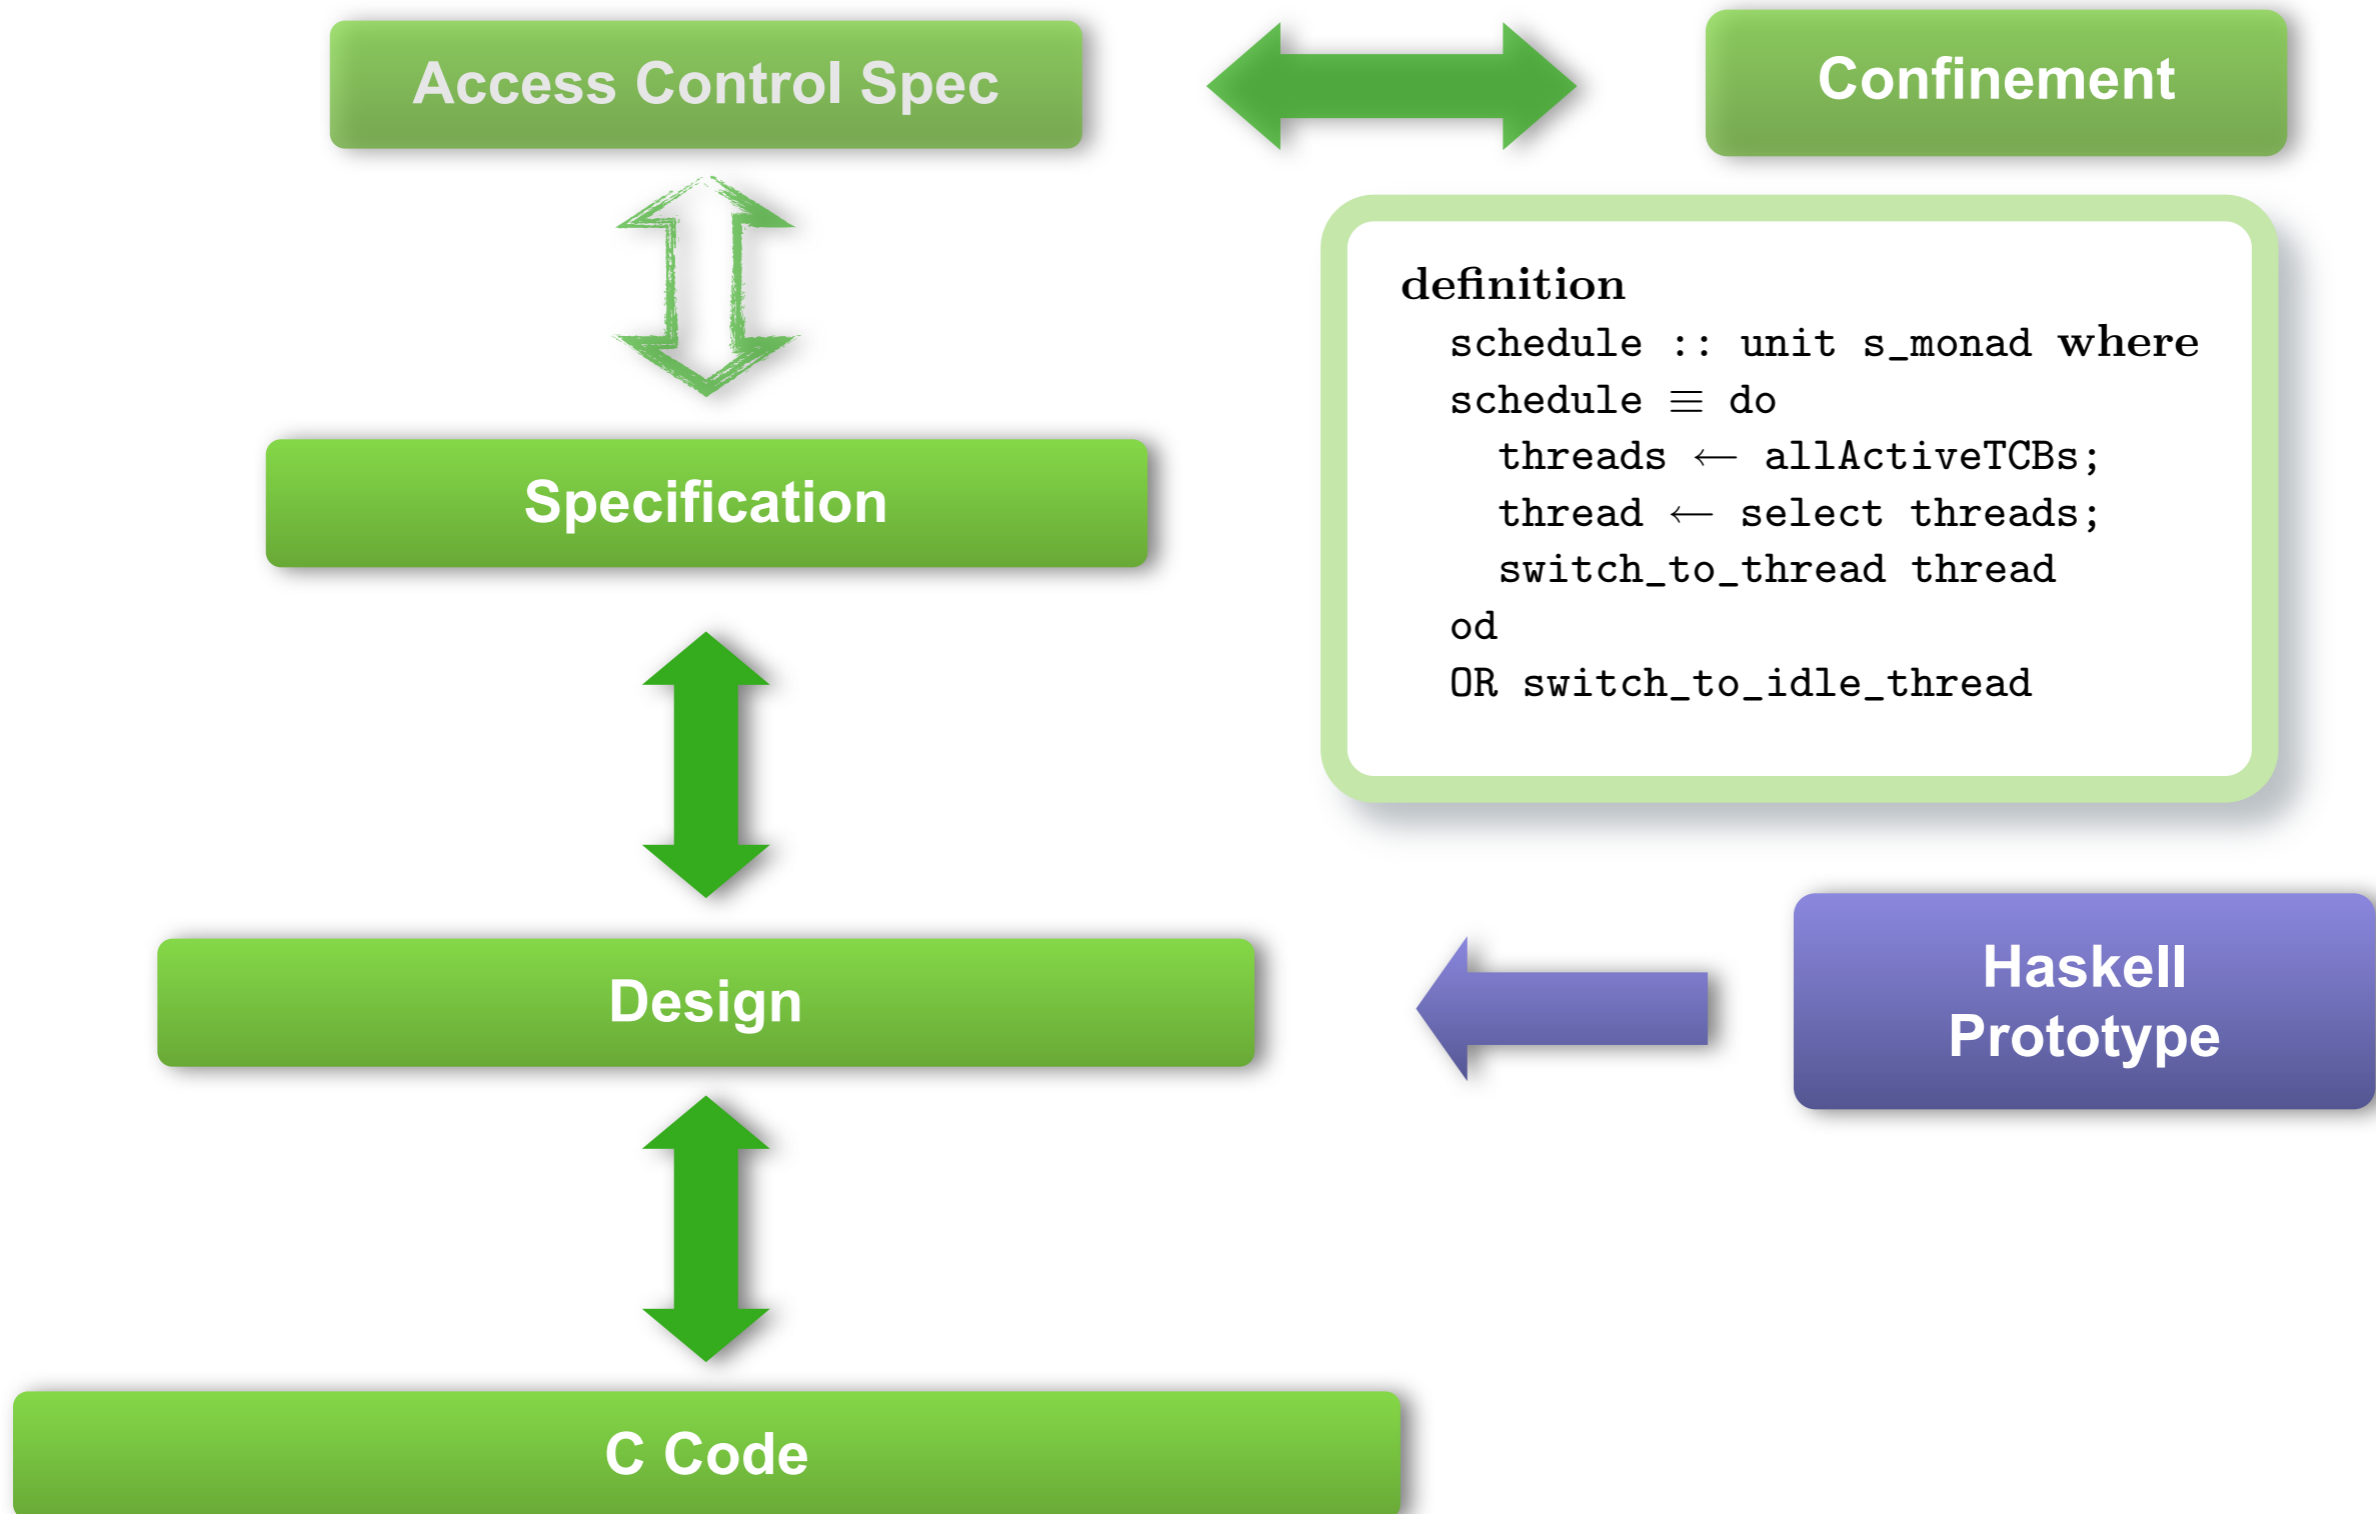
# Proof Architecture



# Proof Architecture



# Proof Architecture



# Proof Architecture

Access Control Spec

Confinement

Specification

Design

C Code

```
schedule :: Kernel ()
schedule = do
  action <- getSchedulerAction
  case action of
    ResumeCurrentThread -> return ()
    ChooseNewThread -> do
      chooseThread
      setSchedulerAction ResumeCurrentThread
    SwitchToThread t -> do
      switchToThread t
      setSchedulerAction ResumeCurrentThread

chooseThread :: Kernel ()
chooseThread = do
  r <- findM chooseThread' (reverse [minBound .. maxBound])
  when (r == Nothing) $ switchToIdleThread
  where
```

# Proof Architecture



Access Control Spec

Confinement

Specification

Design

C Code


```
void
schedule(void) {
    switch ((word_t)ksSchedulerAction) {
        case (word_t)SchedulerAction_ResumeCurrentThread:
            break;

        case (word_t)SchedulerAction_ChooseNewThread:
            chooseThread();
            ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
            break;

        default: /* SwitchToThread */
            switchToThread(ksSchedulerAction);
            ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
            break;
    }
}

void
chooseThread(void) {
    prio_t prio;
    tcb_t *thread, *next;
```







seL4

# Did you find any Bugs?



## Bugs found

- in C: 160
- in design: ~150
- in spec: ~150

**460 bugs**

```
void
schedule(void) {
    switch ((word_t)ksSchedulerAction) {
        case (word_t)SchedulerAction_ResumeCurrentThread:
            break;

        case (word_t)SchedulerAction_ChooseNewThread:
            chooseThread();
            ksSchedulerAction = SchedulerAction_ResumeCurrentThread;
            break;
    }
}

void
chooseThr
prio_
tcb_t

for(p
f

if(!isRunnable(thread)) {
    next = thread->tcbSchedNext;
    tcbSchedDequeue(thread);
}
else {
    switchToThread(thread);
    return;
}
}
}

switchToThread(thread());
}
```

## Effort

<b>Haskell design</b>	2 py
<b>First C impl.</b>	2 weeks
<b>Debugging/Testing</b>	2 months
<b>Kernel verification</b>	12 py
<b>Formal frameworks</b>	10 py
<b>Total</b>	25 py




# Did you find any Bugs?



```
void  
schedule(void) {  
    switch ((word_t)ksSchedulerAction) {  
        case (word_t)SchedulerAction_ResumeCurrentThread:  
            break;  
  
        case (word_t)SchedulerAction_ChooseNewThread:  
            chooseThread();  
            ksSchedulerAction = SchedulerAction_ResumeCurrentThread;  
            break;  
    }  
}  
  
void  
chooseThr  
prio_  
tcb_t  
  
for(p  
f  
  
if(!isRunnable(thread)) {  
    next = thread->tcbSchedNext;  
    tcbSchedDequeue(thread);  
}  
else {  
    switchToThread(thread);  
    return;  
}  
}
```

## Bugs found

- in C: 160
  - in design: ~150
  - in spec: ~150
- 460 bugs**



### Effort

Haskell design	2 py
First C impl.	2 weeks
Debugging/Testing	2 months
Kernel verification	12 py
Formal frameworks	10 py
<b>Total</b>	<b>25 py</b>

# Access Control




# Access Control





# Proof Architecture




# Take-Grant model

## Lipton and Snyder:

- entities represented as nodes of a graph
- capabilities represented as edges of a graph
- rights are contained in capabilities


**The Rights:**

- Read
- Write
- Create
- Take
- Grant




# Operations - Create

---



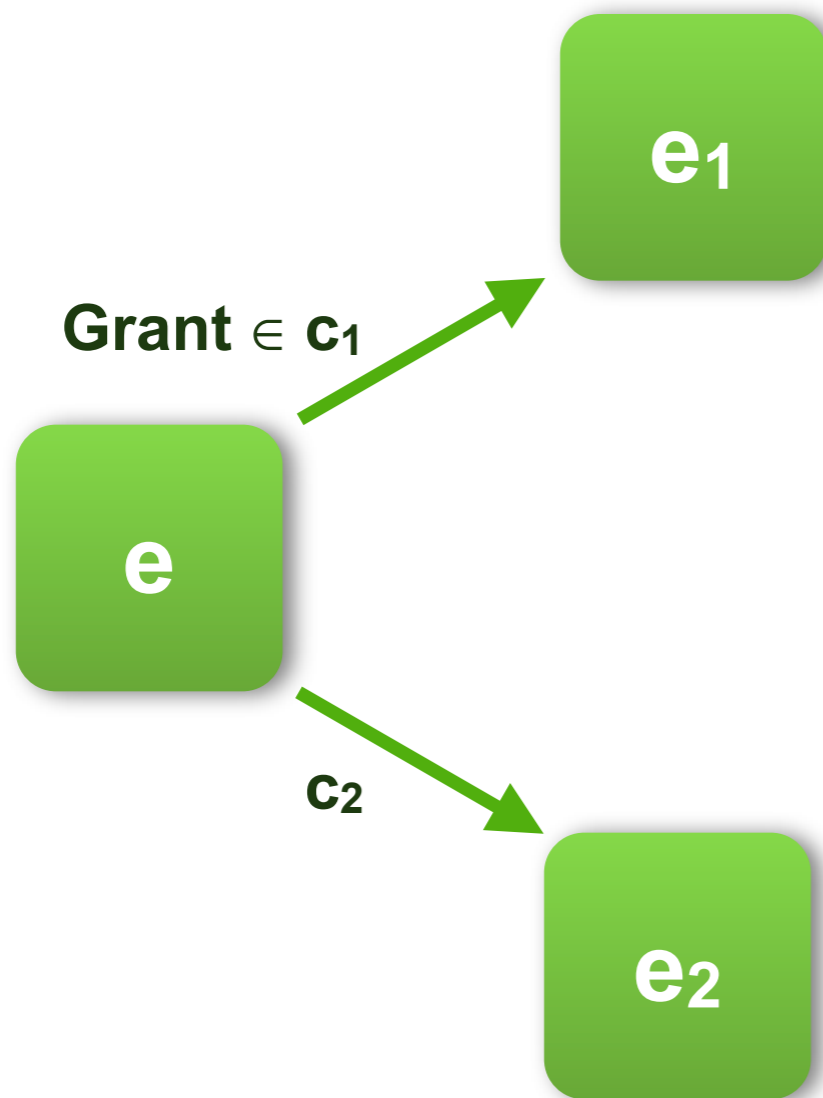
**Create new entity**

# Operations - Create



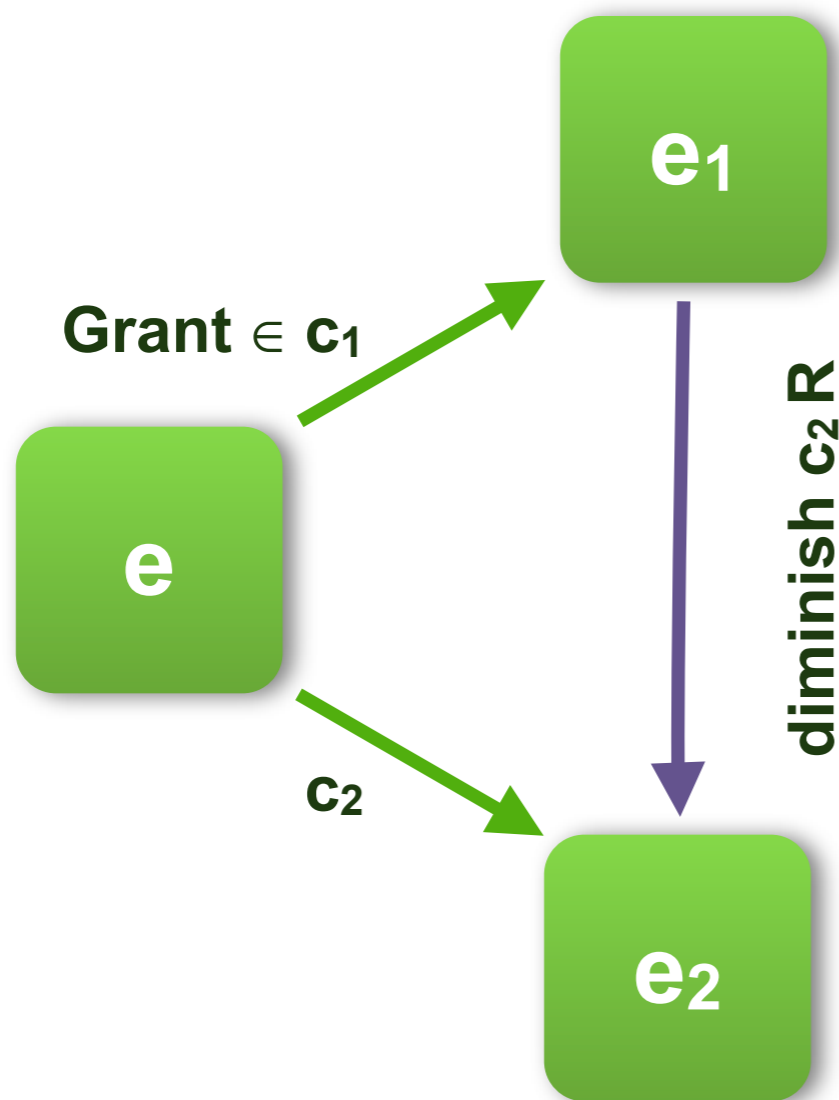
**Create new entity**

# Operations - Grant



**Grant  $c_2$  to  $e_1$**   
with mask  $R$

# Operations - Grant



**Grant**  $c_2$  to  $e_1$   
with mask  $R$




# Operations - Remove/Delete

## Remove capability $c_2$



## Delete entity $e_2$



# Operations - Remove/Delete

## Remove capability $c_2$



## Delete entity $e_2$



# Operations - Remove/Delete

## Remove capability $c_2$



## Delete entity $e_2$



# Questions

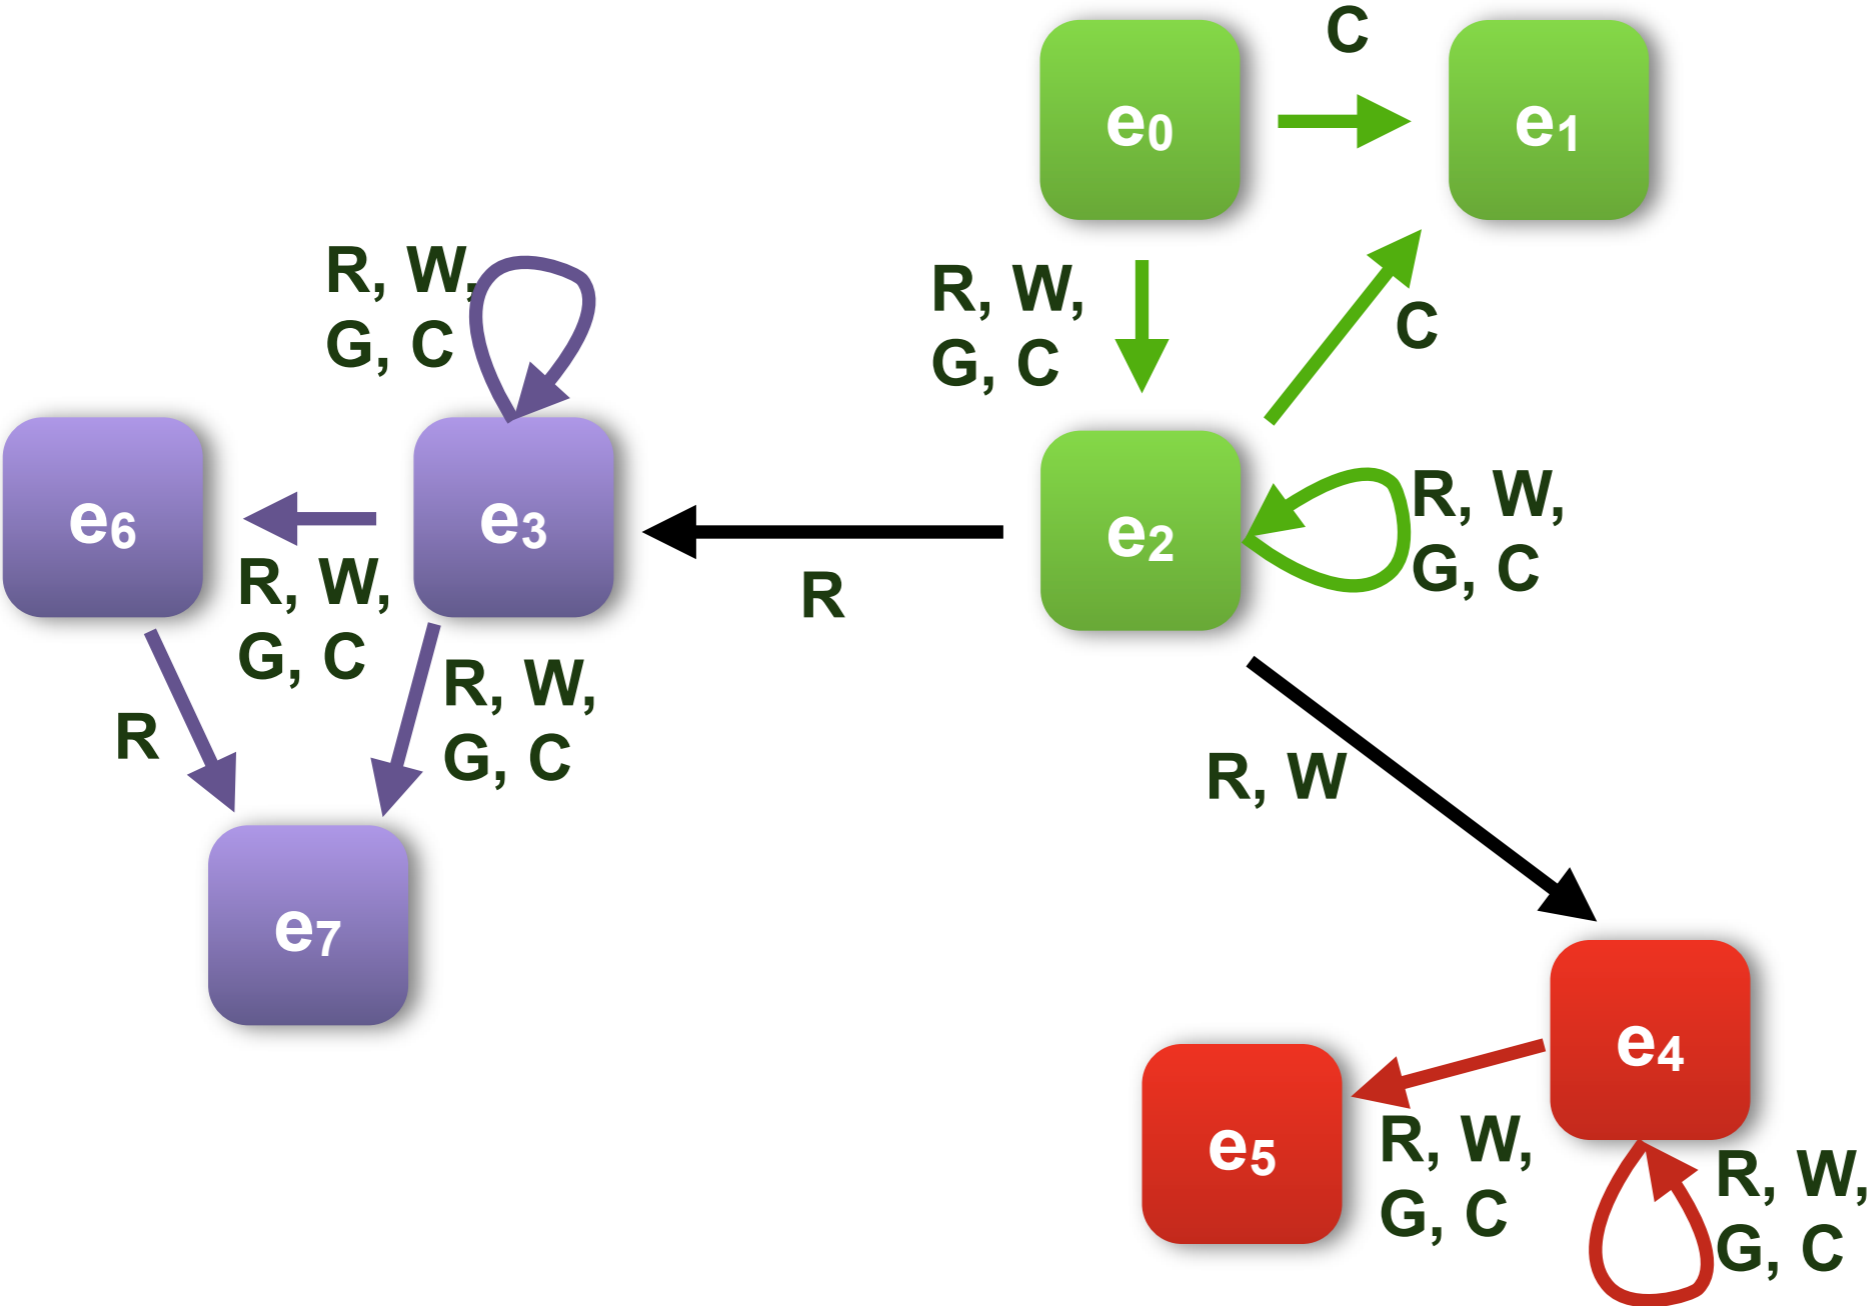
---

## For any state in the future:


- Can **E** gain authority to do **X**?
- Can **E** gain more **authority** than it has?
- **How** much more?
- Can **information** flow from A to B?



# Example

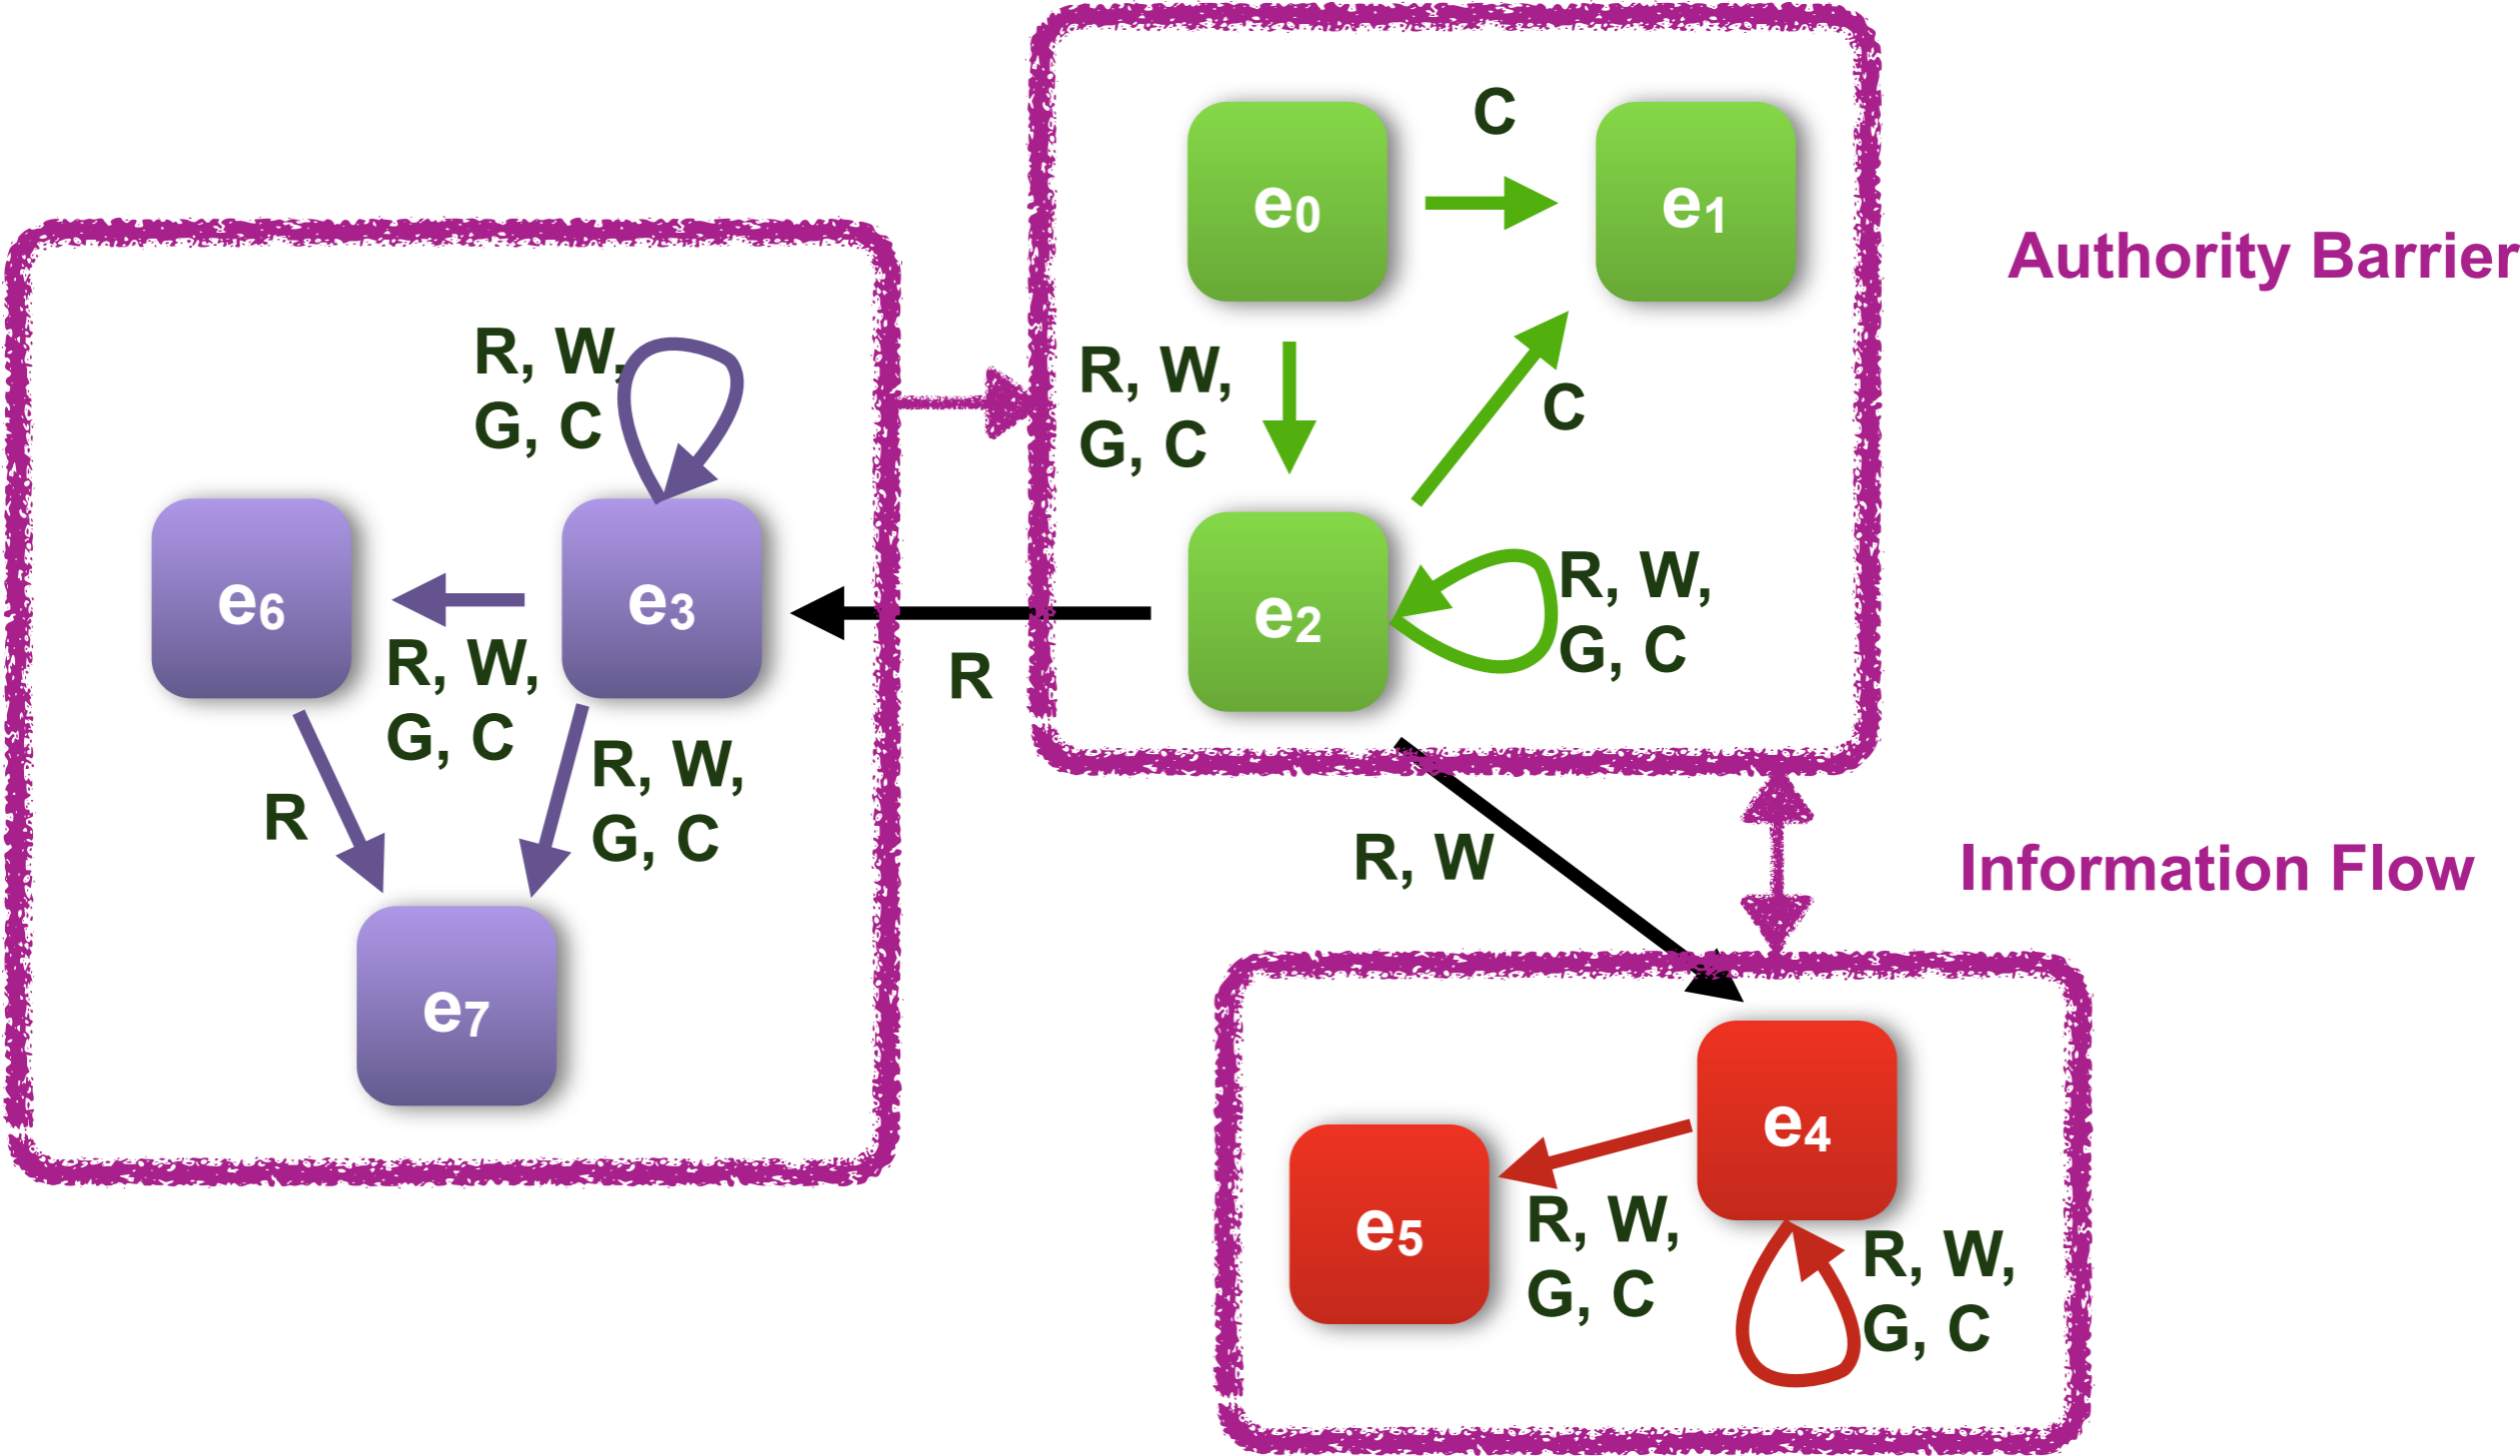


# Example



Authority Barrier

# Example








# Now What?

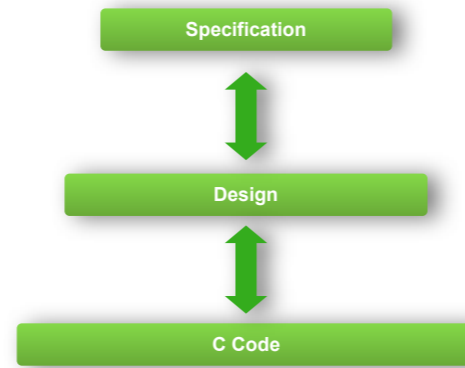


# Current Proof



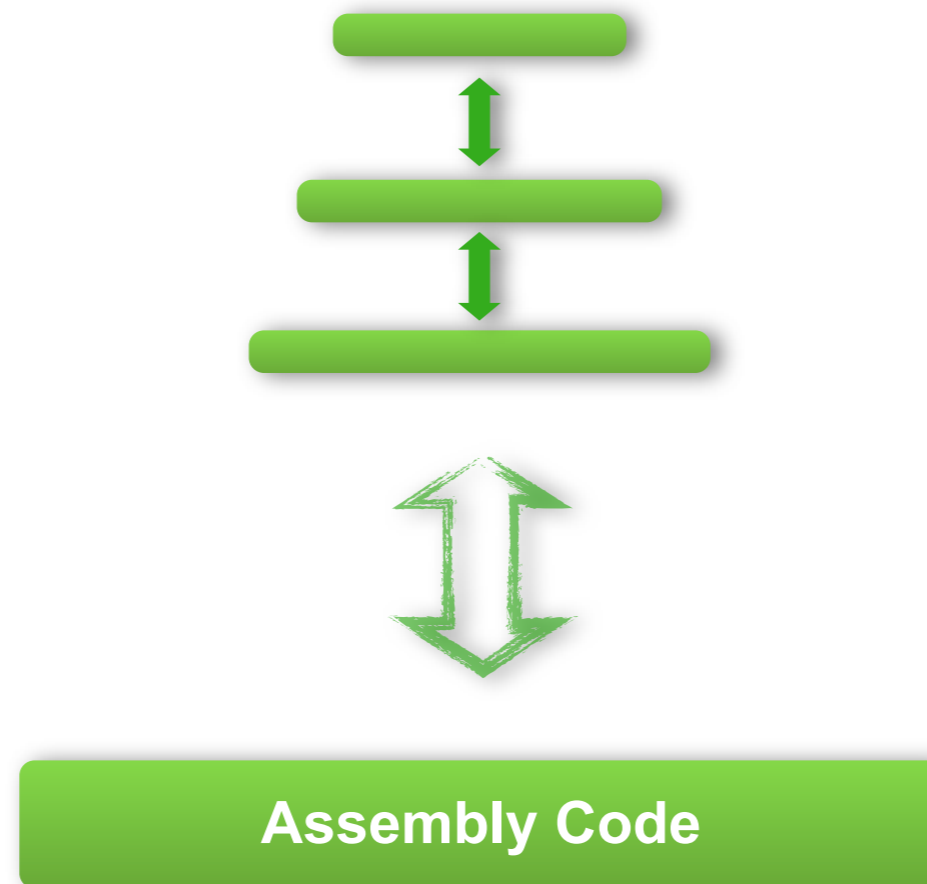
# Current Proof

---




# Even More Assurance?

---





# Even More Assurance?




Assembly Code

## Assume correct:

- compiler + linker (wrt. C opsem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

# Even More Assurance?

**Compiler Verification**  
CompCert

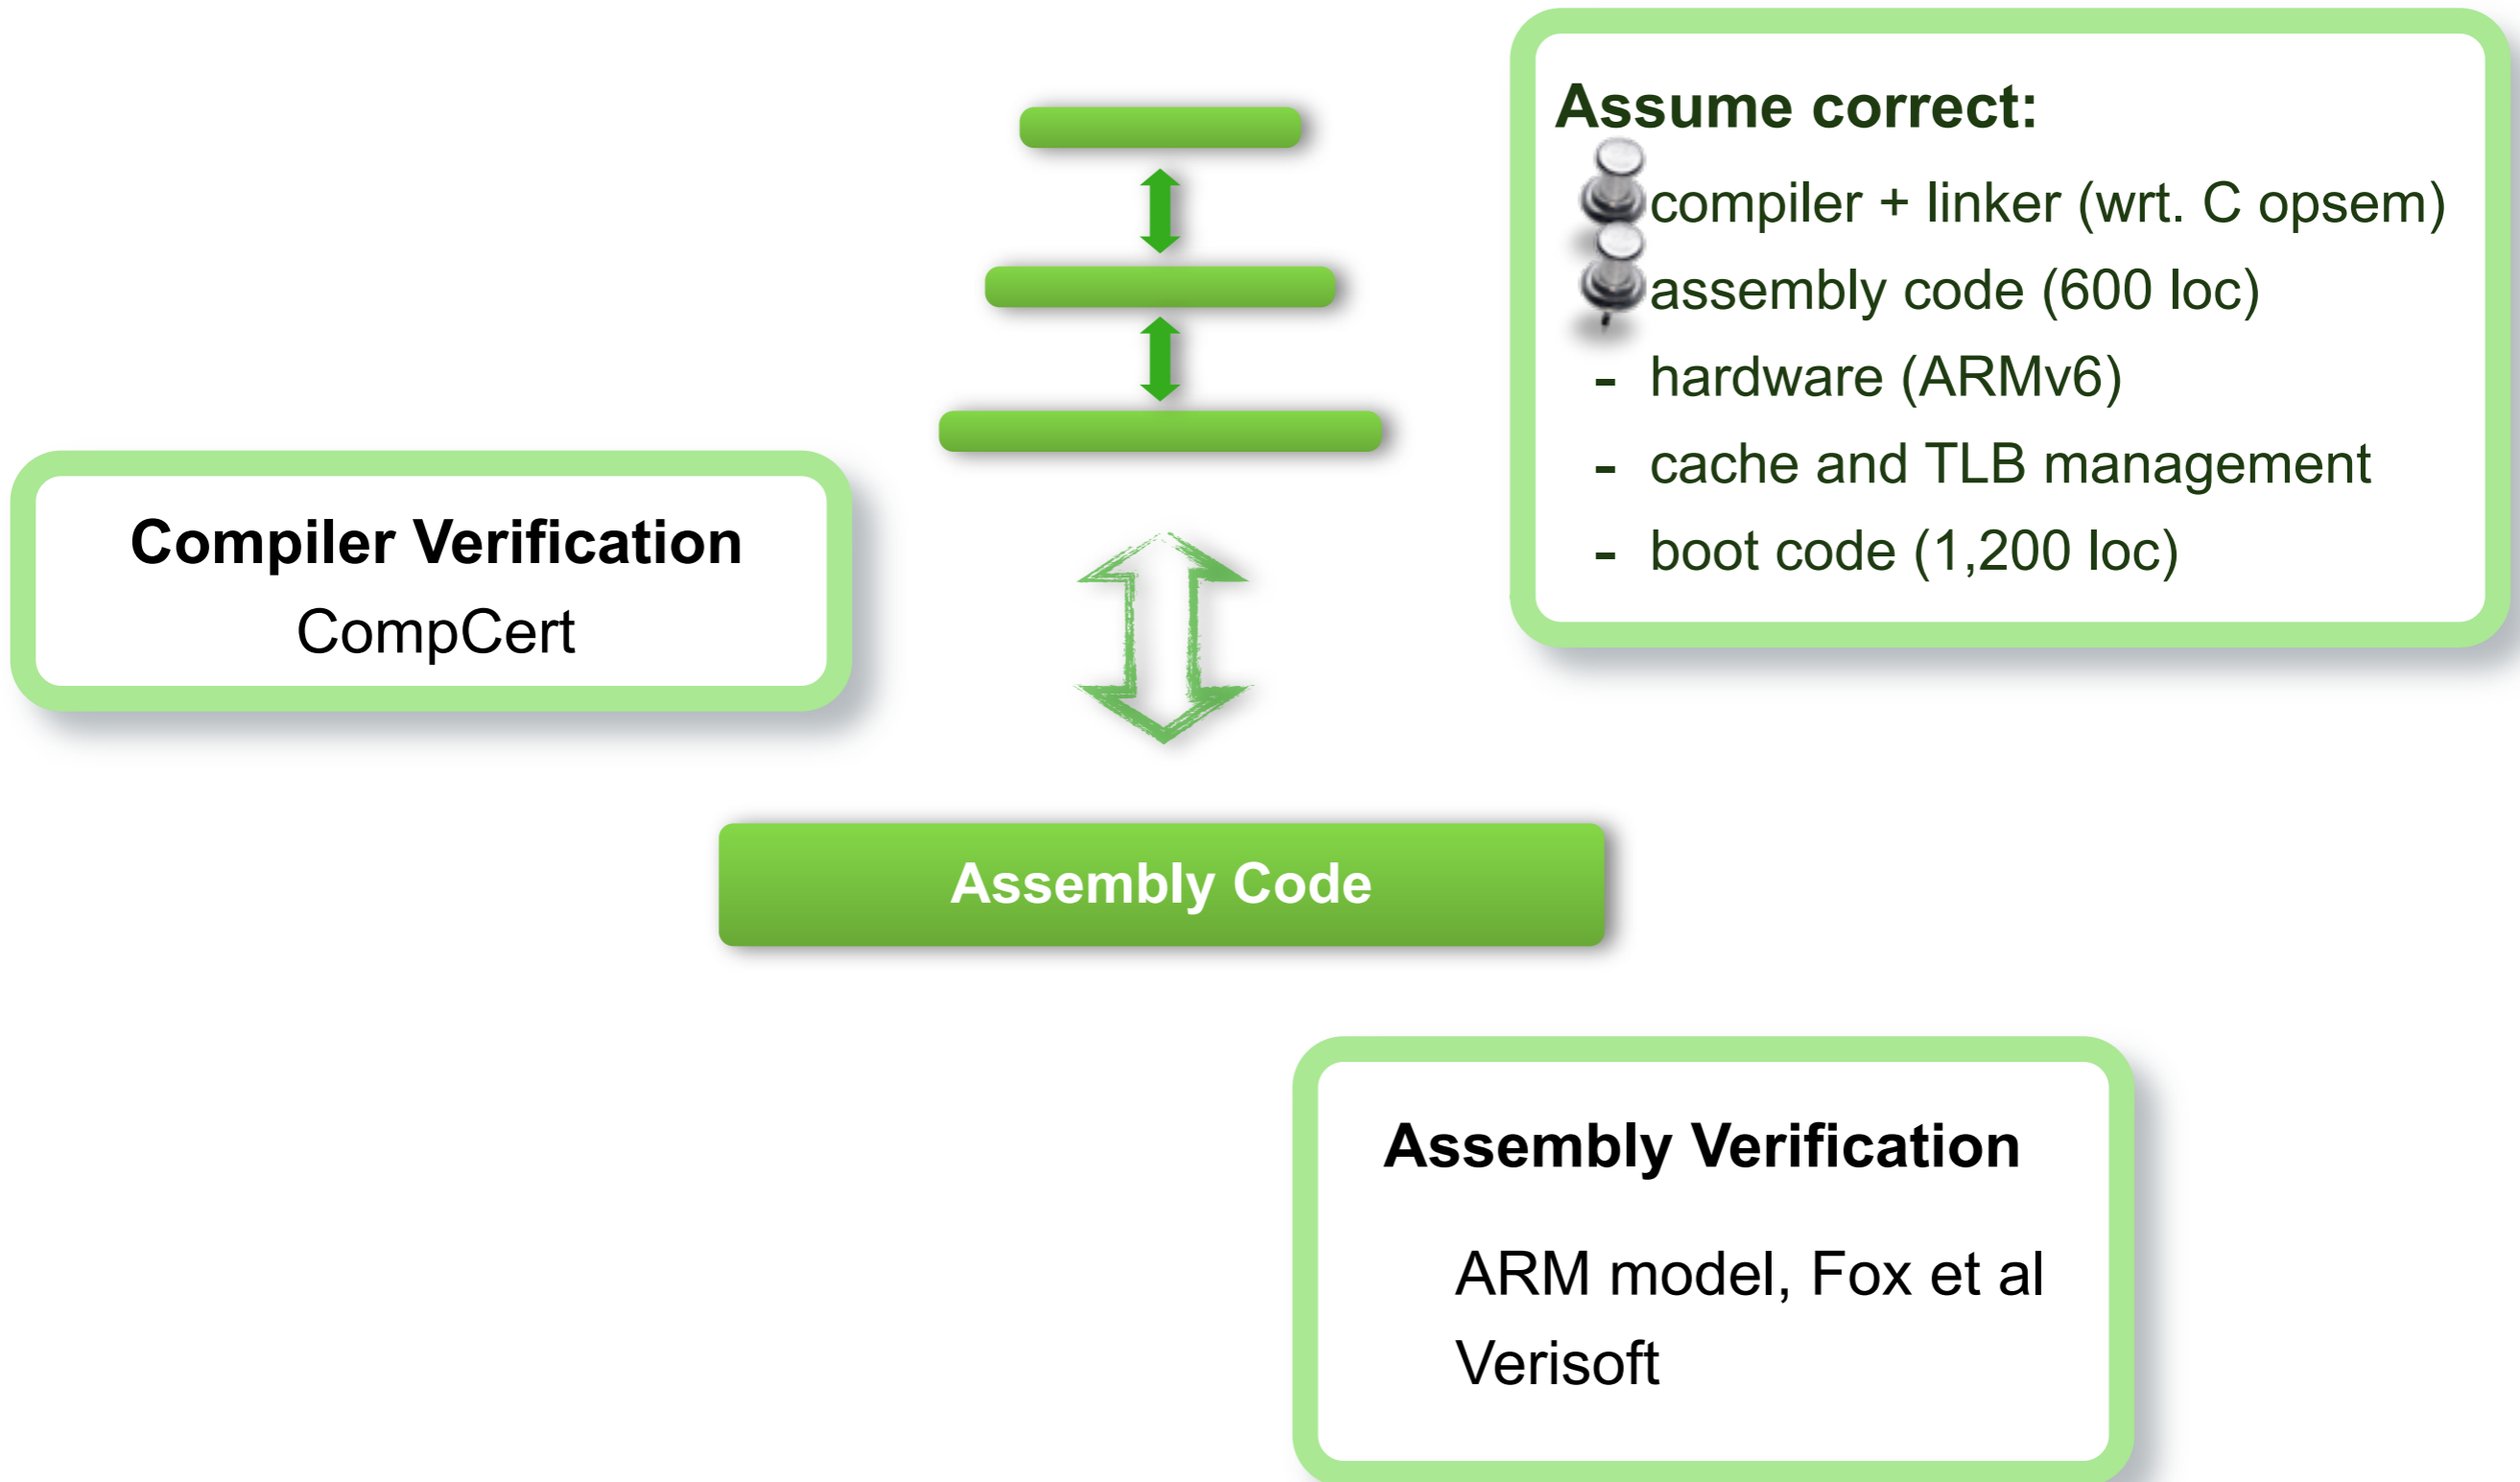


**Assembly Code**

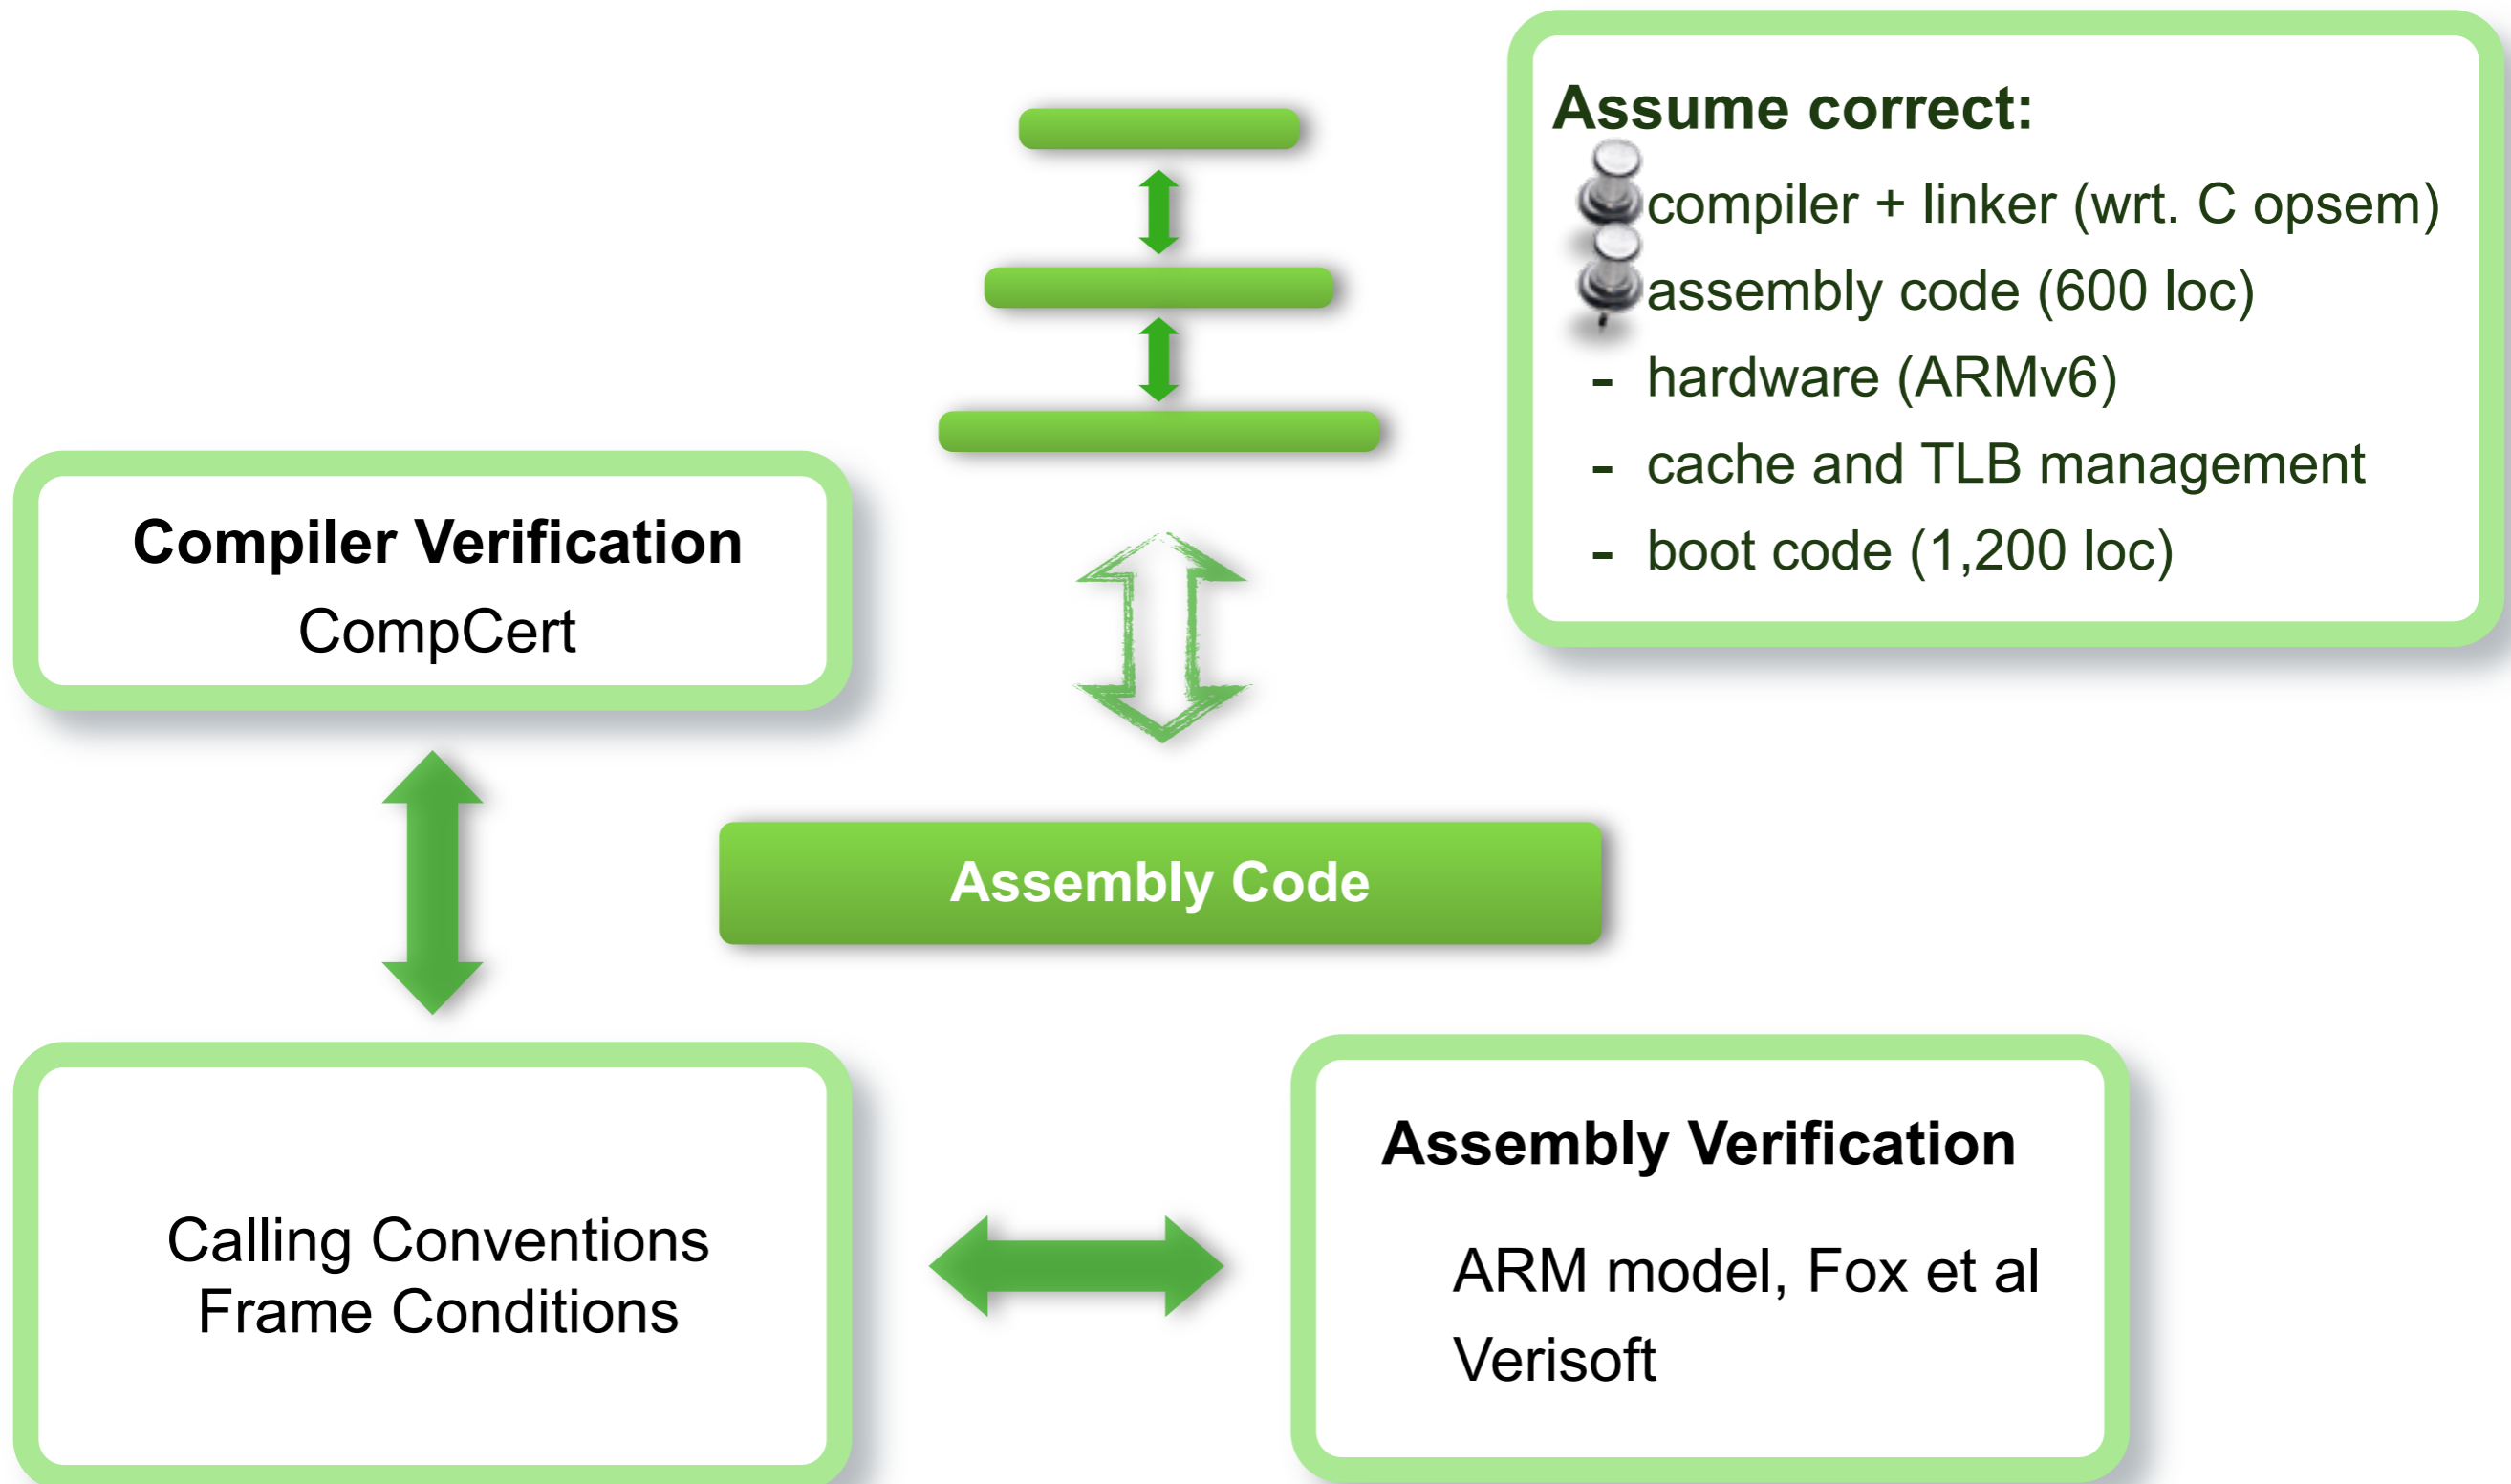
## Assume correct:

- 🔍 compiler + linker (wrt. C opsem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

# Even More Assurance?

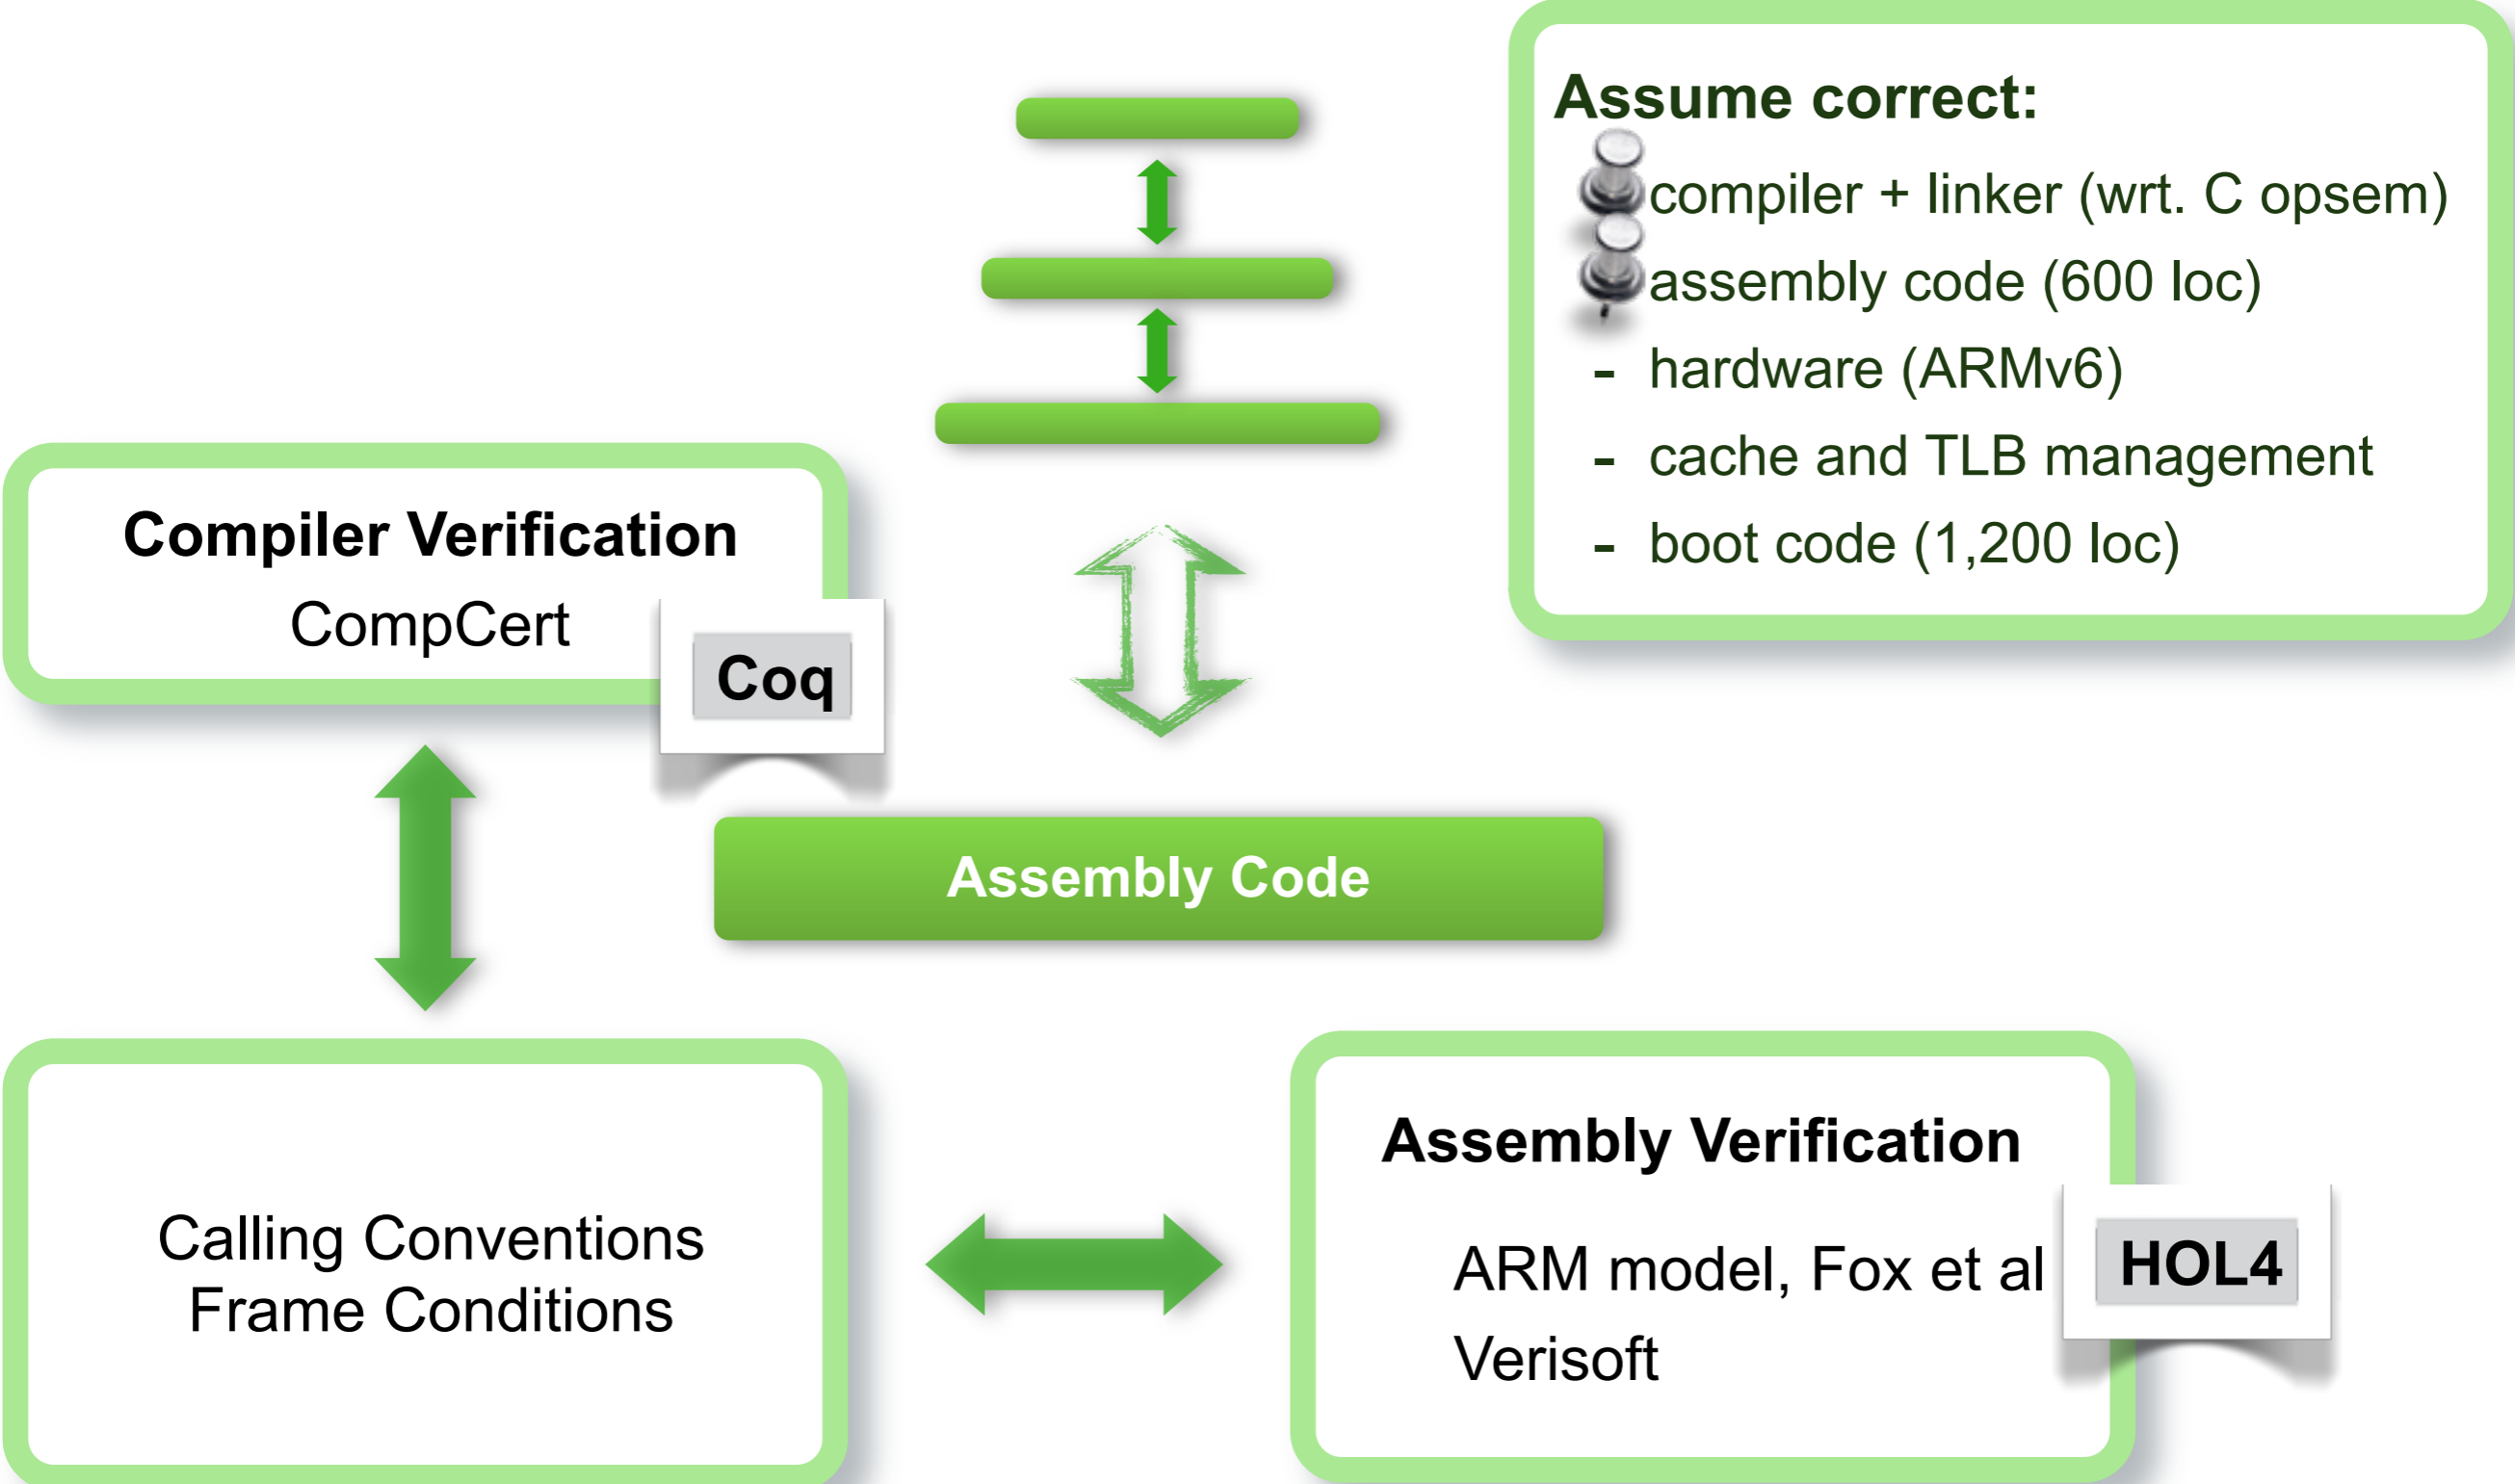


# Even More Assurance?

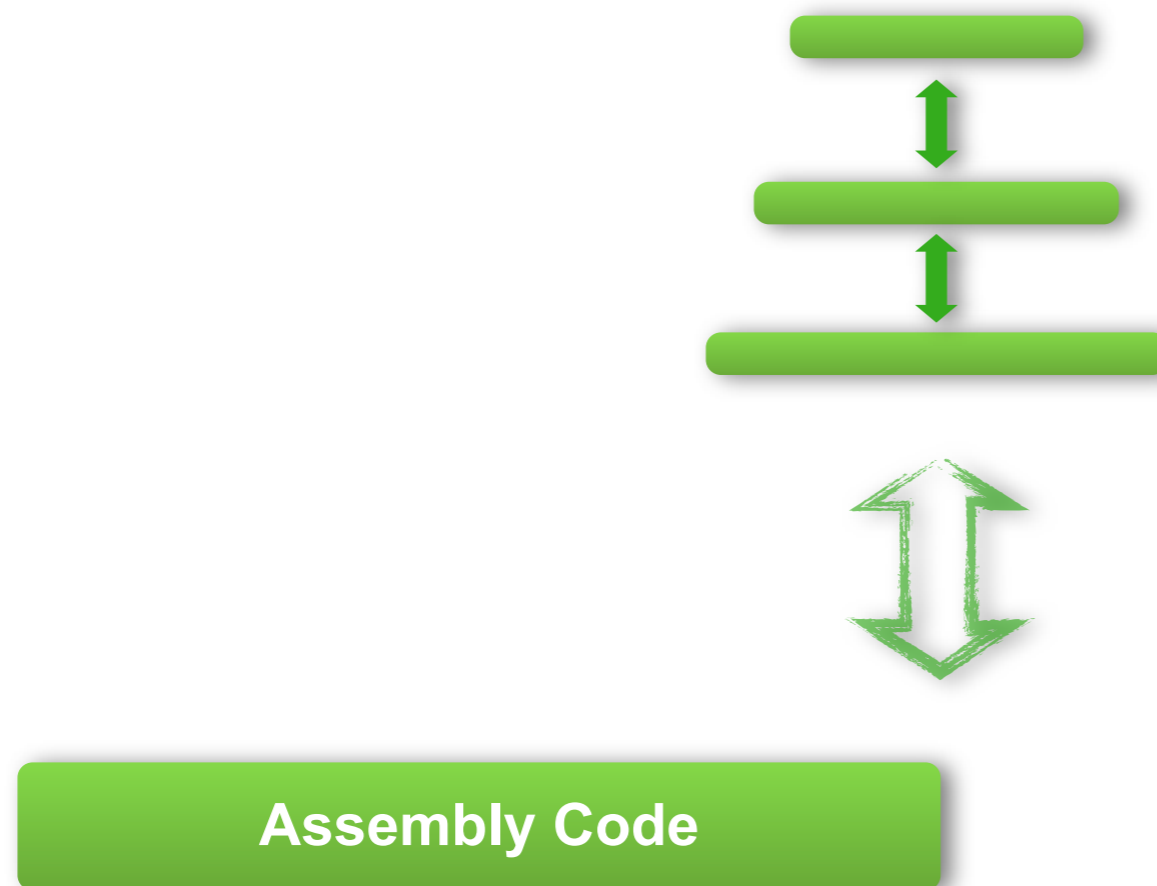




# Even More Assurance?




# Even More Assurance?



## Assume correct:

- compiler + linker (wrt. C opsem)
- assembly code (600 loc)
  - hardware (ARMv6)
  - cache and TLB management
  - boot code (1,200 loc)

# Even More Assurance?



## Assume correct:


- compiler + linker (wrt. C opsem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

Assembly Code

Cache/TLB model

Verisoft XT

# Even More Assurance?



**Assume correct:**

- compiler + linker (wrt. C opsem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

Assembly Code


Cache/TLB model

Verisoft XT



# Even More Assurance?

**boot result verification**



## Assume correct:

- compiler + linker (wrt. C opsem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

Assembly Code

Cache/TLB model

Verisoft XT



**Hardware Verification**


VHDL?

VAMP  
FM9001  
Intel i7

# Even More Assurance?



**boot result verification**



**Assume correct:**

- compiler + linker (wrt. C opsem)
- assembly code (600 loc)
- hardware (ARMv6)
- cache and TLB management
- boot code (1,200 loc)

**Assembly Code**

**Cache/TLB model**

Verisoft XT




**Hardware Verification**

**VHDL?**

VAMP  
FM9001  
Intel i7

# Systems On Top

---



# Systems On Top

---



# Systems On Top


**C/C++  
progs**

**seL4/  
Linux**

**C system call  
bindings**

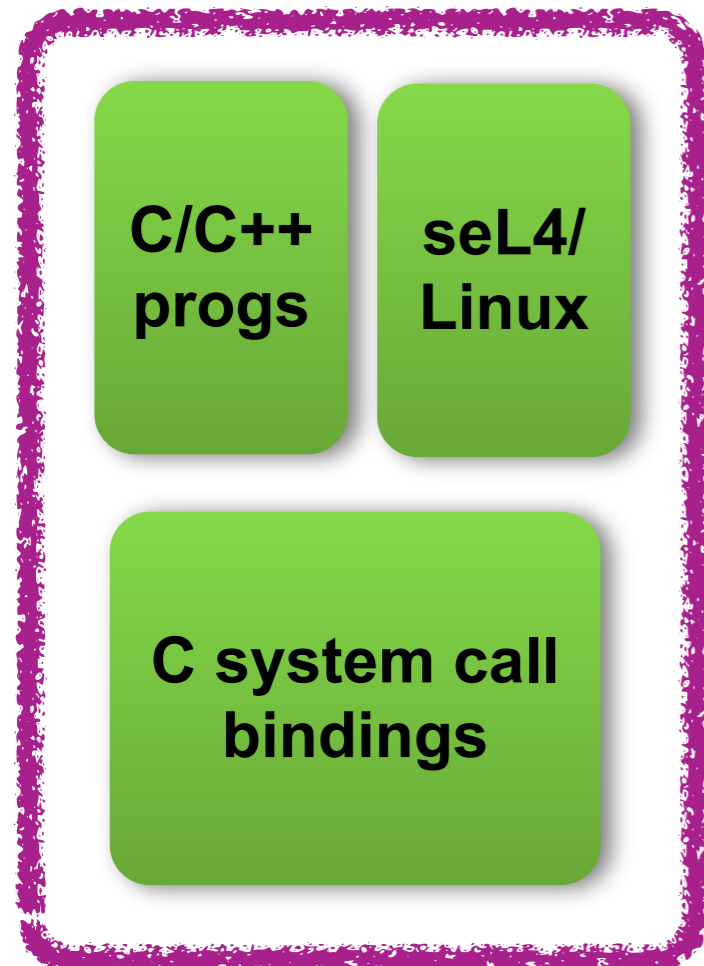
## **Exists:**

- standard seL4 library
- used in seL4/Linux
- not hard to formally verify
- verification scheduled






# Systems On Top

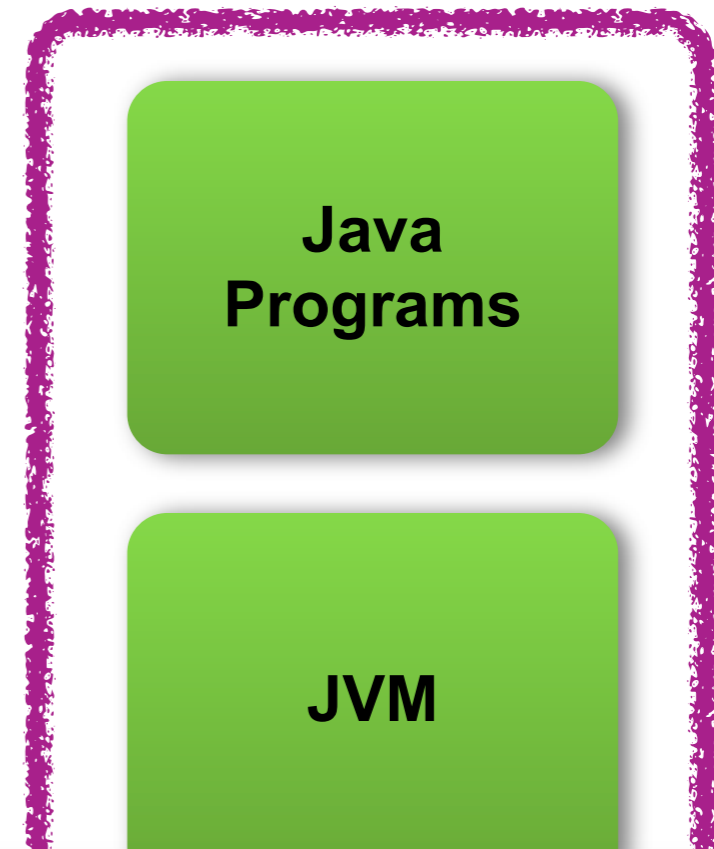
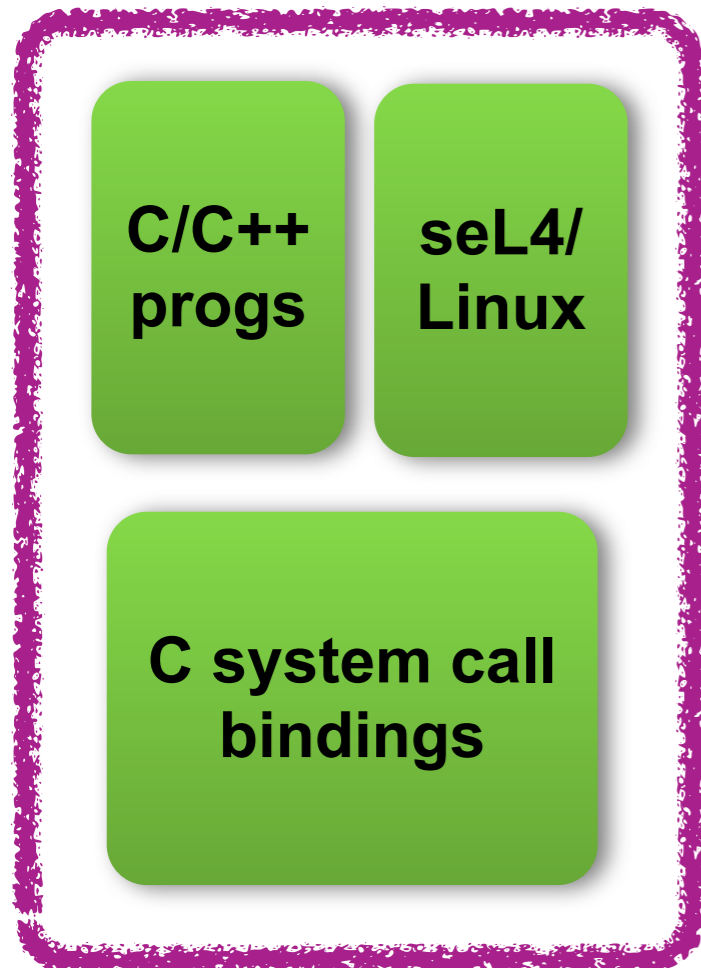


## seL4/Haskell:

- early prototype Haskell runtime
- has seL4 systems call bindings
- verification hard
- runtime verification progress in HASP project @ PSU & Galois




# Systems On Top

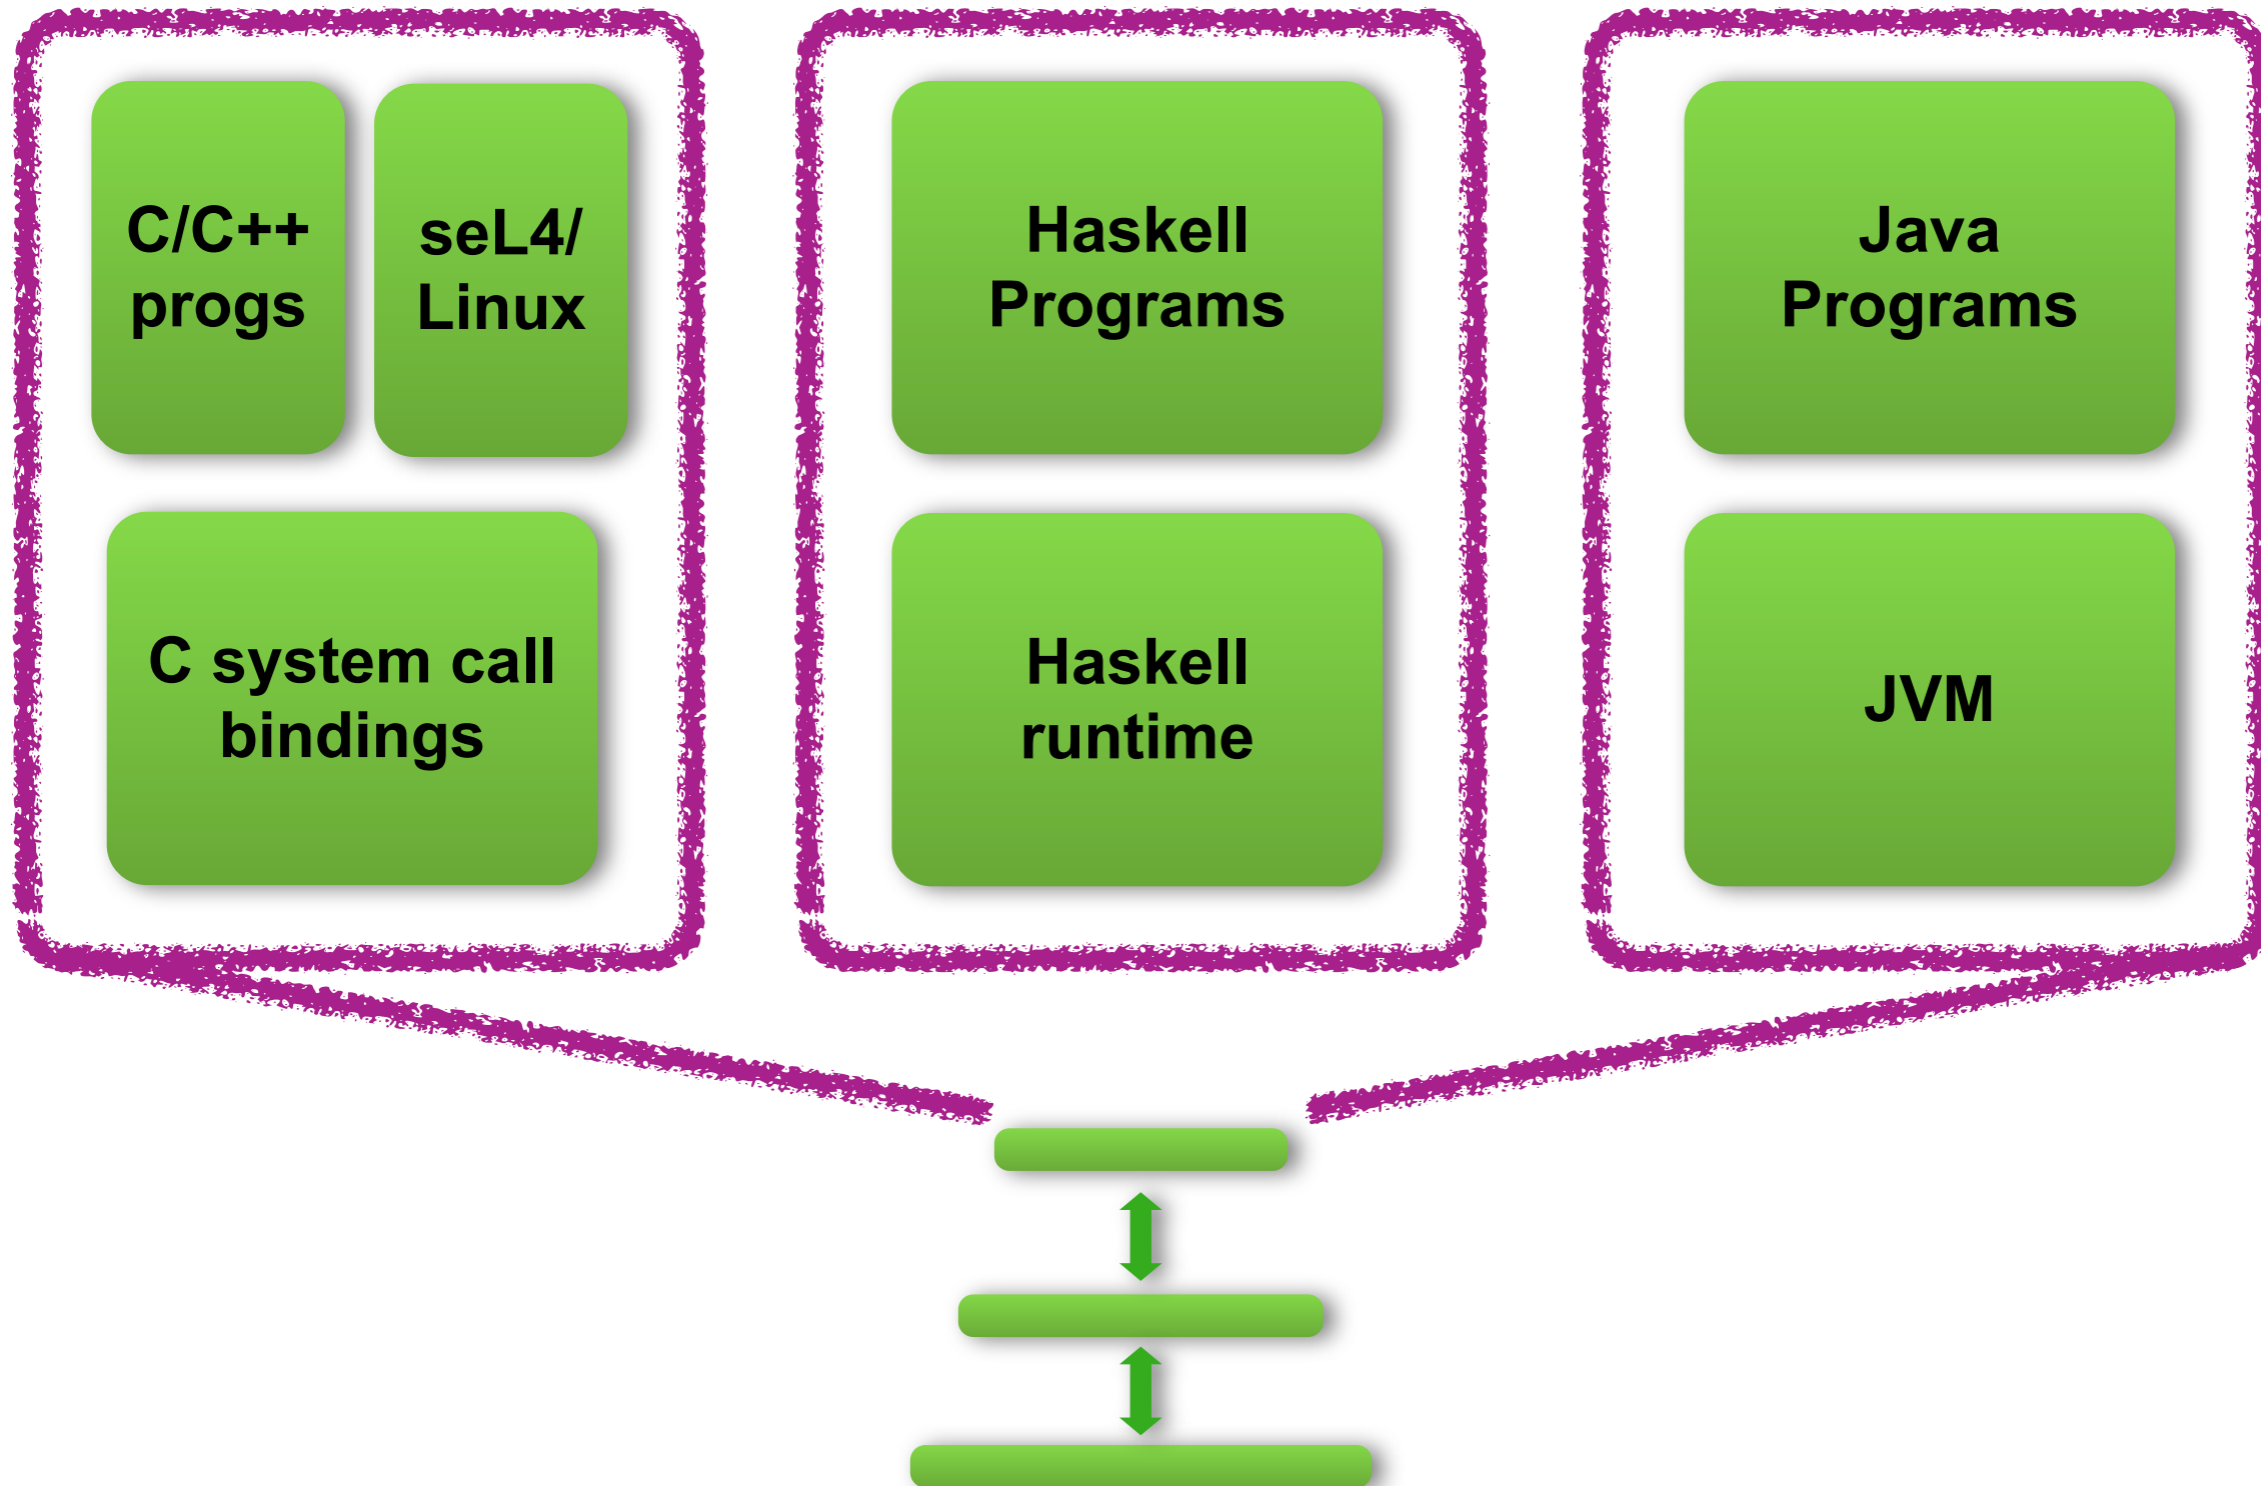


## seL4/JVM:

- any takers?
- JVM extensively formalised
- widely used
- EAL7 smart card implementations exist




# Systems On Top



# Other Architectures

---




# Other Architectures

## seL4/x86:

- x86 version exists, supports Linux
- verification likely, not started yet

## Intel 32bit






# Other Architectures

## seL4/x86:

- x86 version exists, supports Linux
- verification likely, not started yet
  
- Intel VT-d/IOMMU implemented
- enables untrusted device DMA
- verification possible

## Intel 32bit + IOMMU




# Other Architectures

## Intel 32bit + IOMMU + multi core

### seL4/x86:

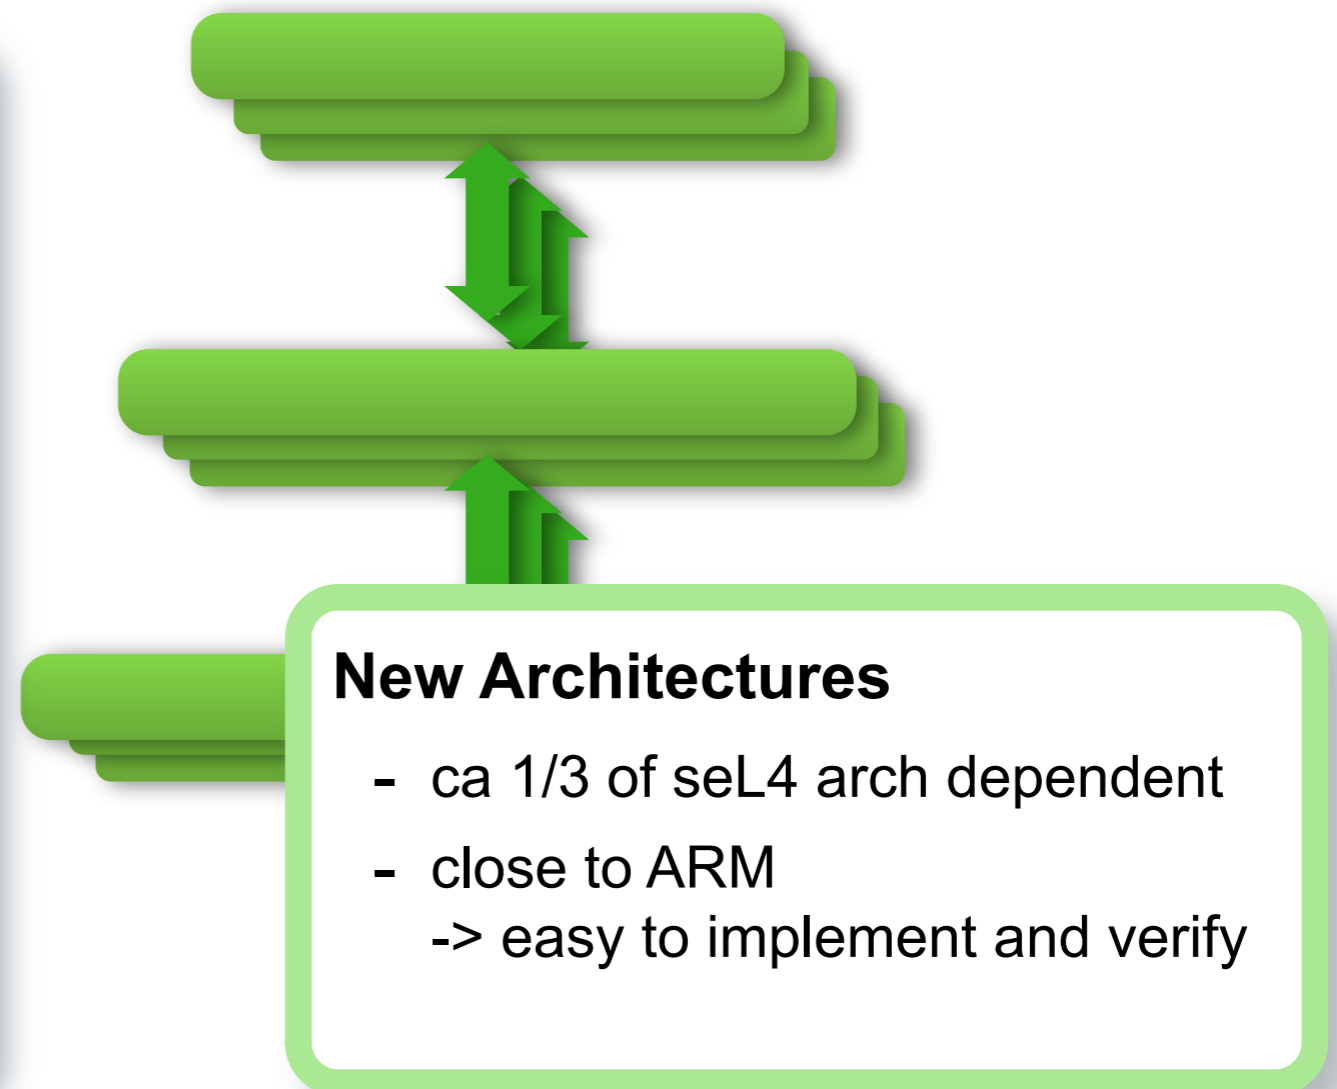
- x86 version exists, supports Linux
- verification likely, not started yet
- Intel VT-d/IOMMU implemented
- enables untrusted device DMA
- verification possible
- experimental multi processor version
- initial proofs exist



## Intel 32bit + IOMMU + multi core

### seL4/x86:

- x86 version exists, supports Linux
- verification likely, not started yet
- Intel VT-d/IOMMU implemented
- enables untrusted device DMA
- verification possible
- experimental multi processor version
- initial proofs exist




### New Architectures

- ca 1/3 of seL4 arch dependent
- close to ARM
- > easy to implement and verify

**Looking Forward**



# Looking Forward





# Trustworthy Embedded Systems

- **L4.verified:**  
functional correctness  
10,000 loc




- **Next step:**  
formal guarantees for  
**> 1,000,000 loc**



# How?


## Exploit:

- seL4 isolation
- verified properties
- MILS architectures



# Challenges


- Find right architecture
- Security analysis
  - identify trusted components
  - ideally take-grant style
  - behaviour of trusted components
- Code-level theorem in the end
  - connect to kernel proof
  - ideally prove trusted component only



# Example System

---


- Scenario:





# Example System

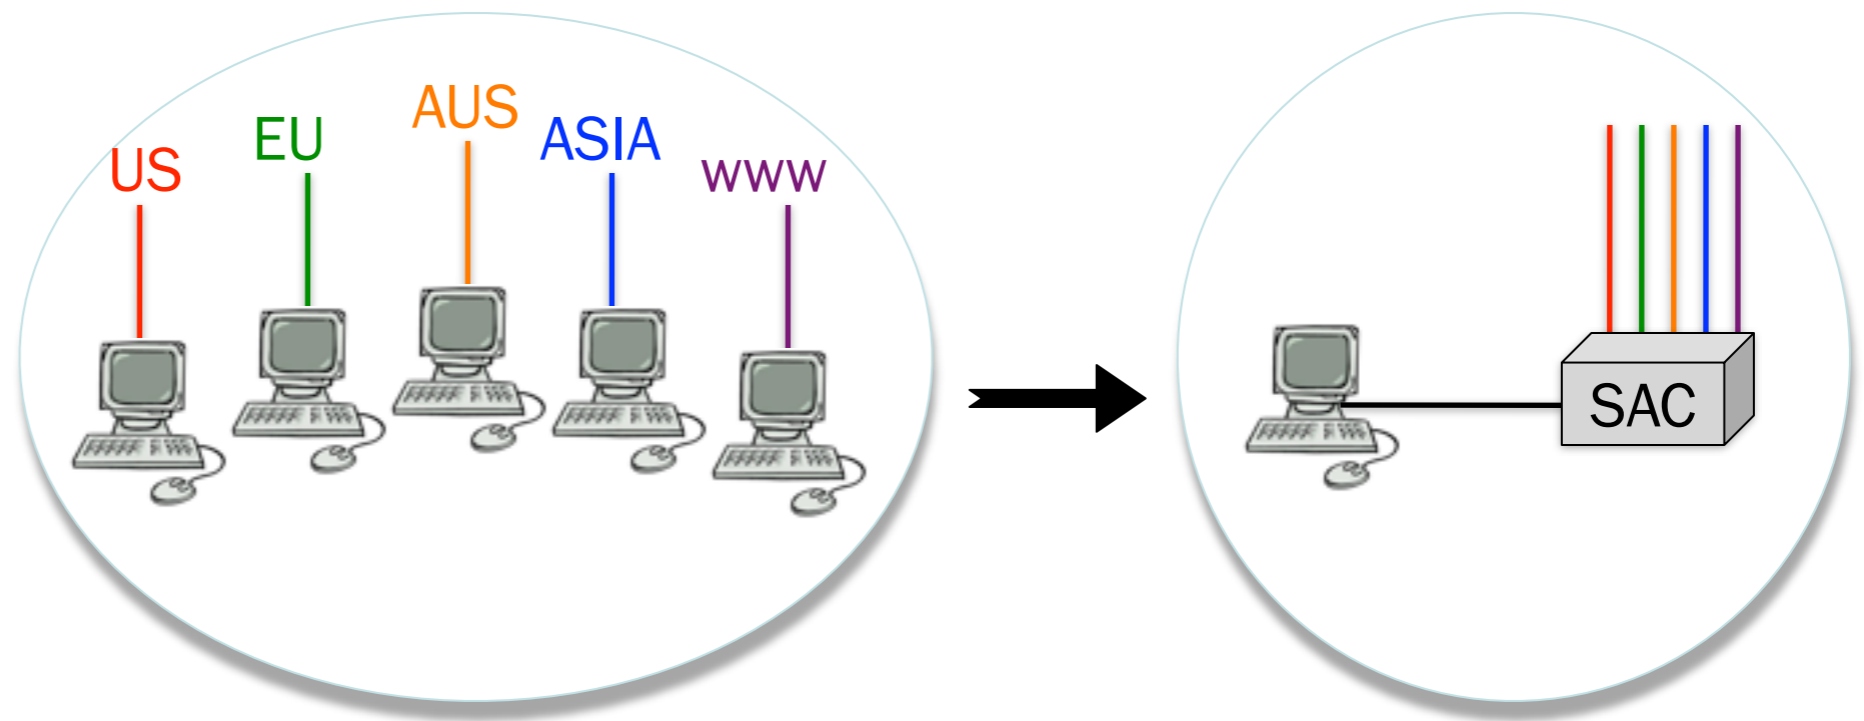
- Scenario:





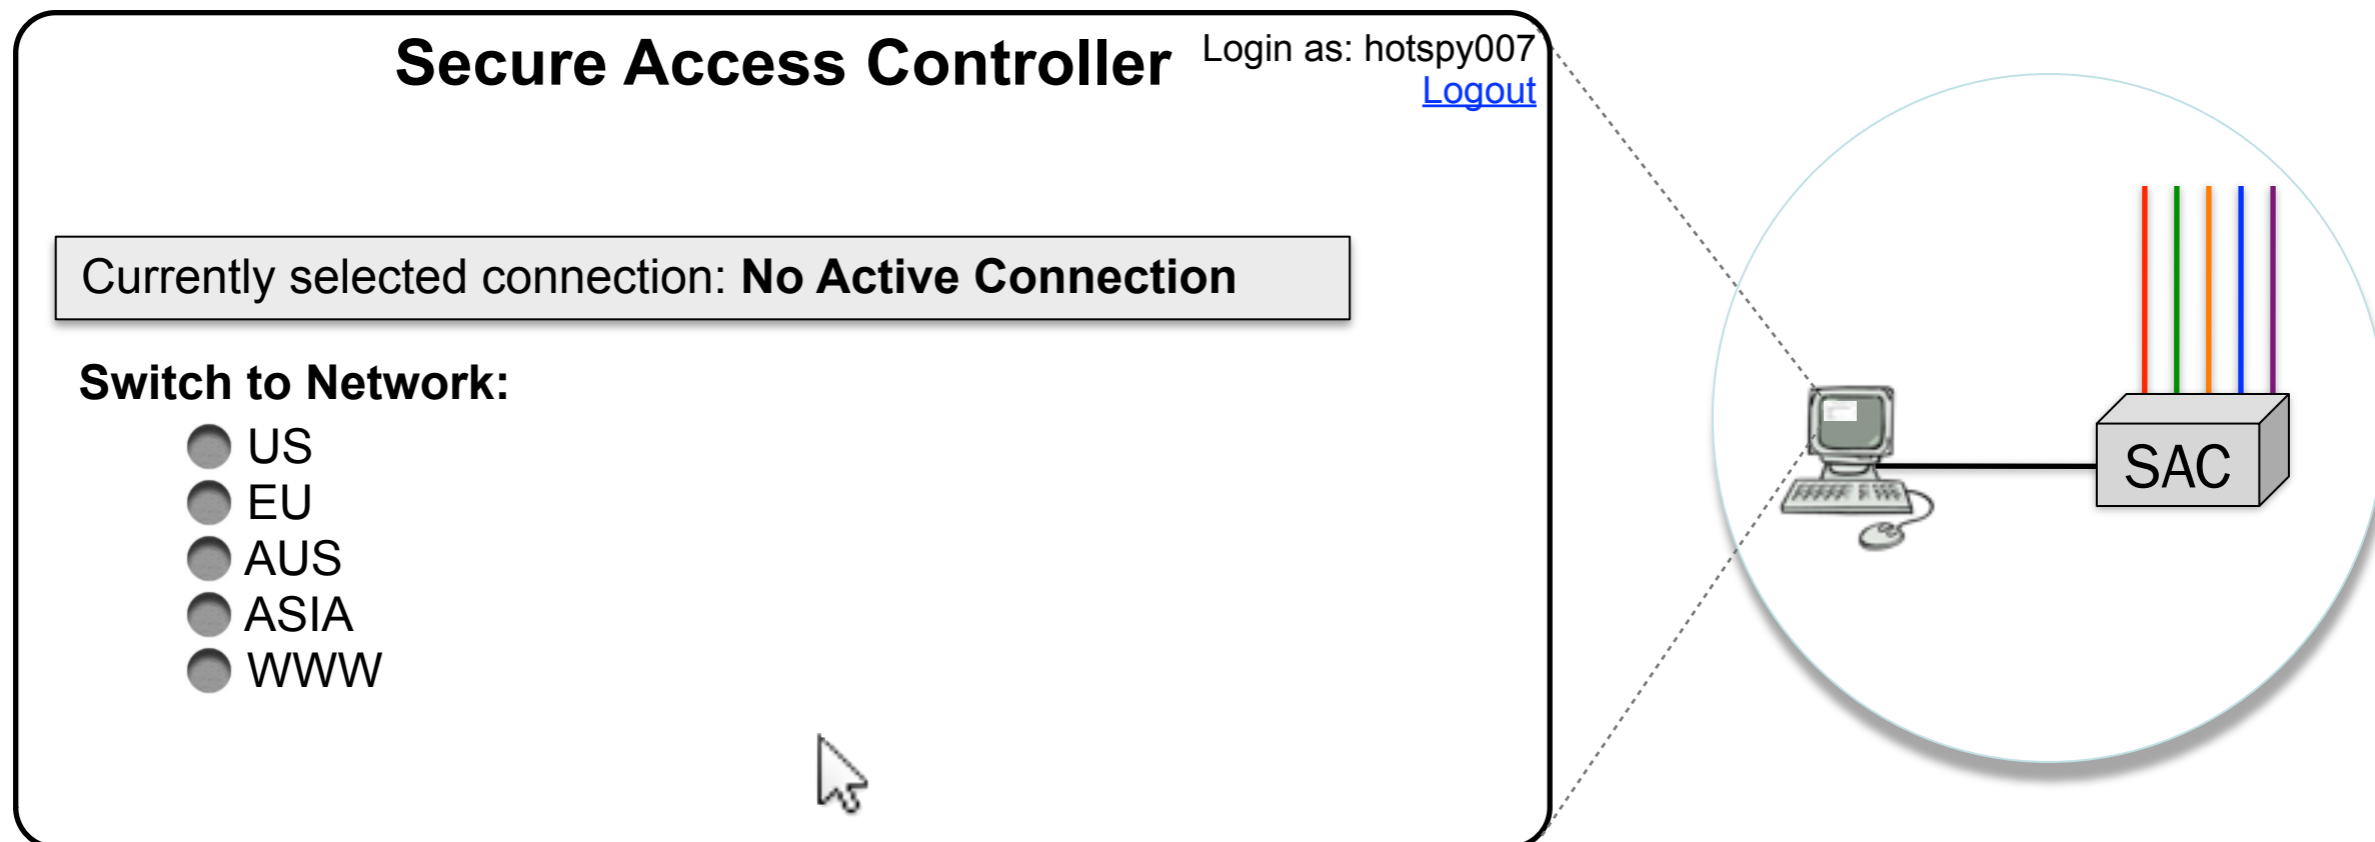
# Example System

- Multilevel Secure Access Device



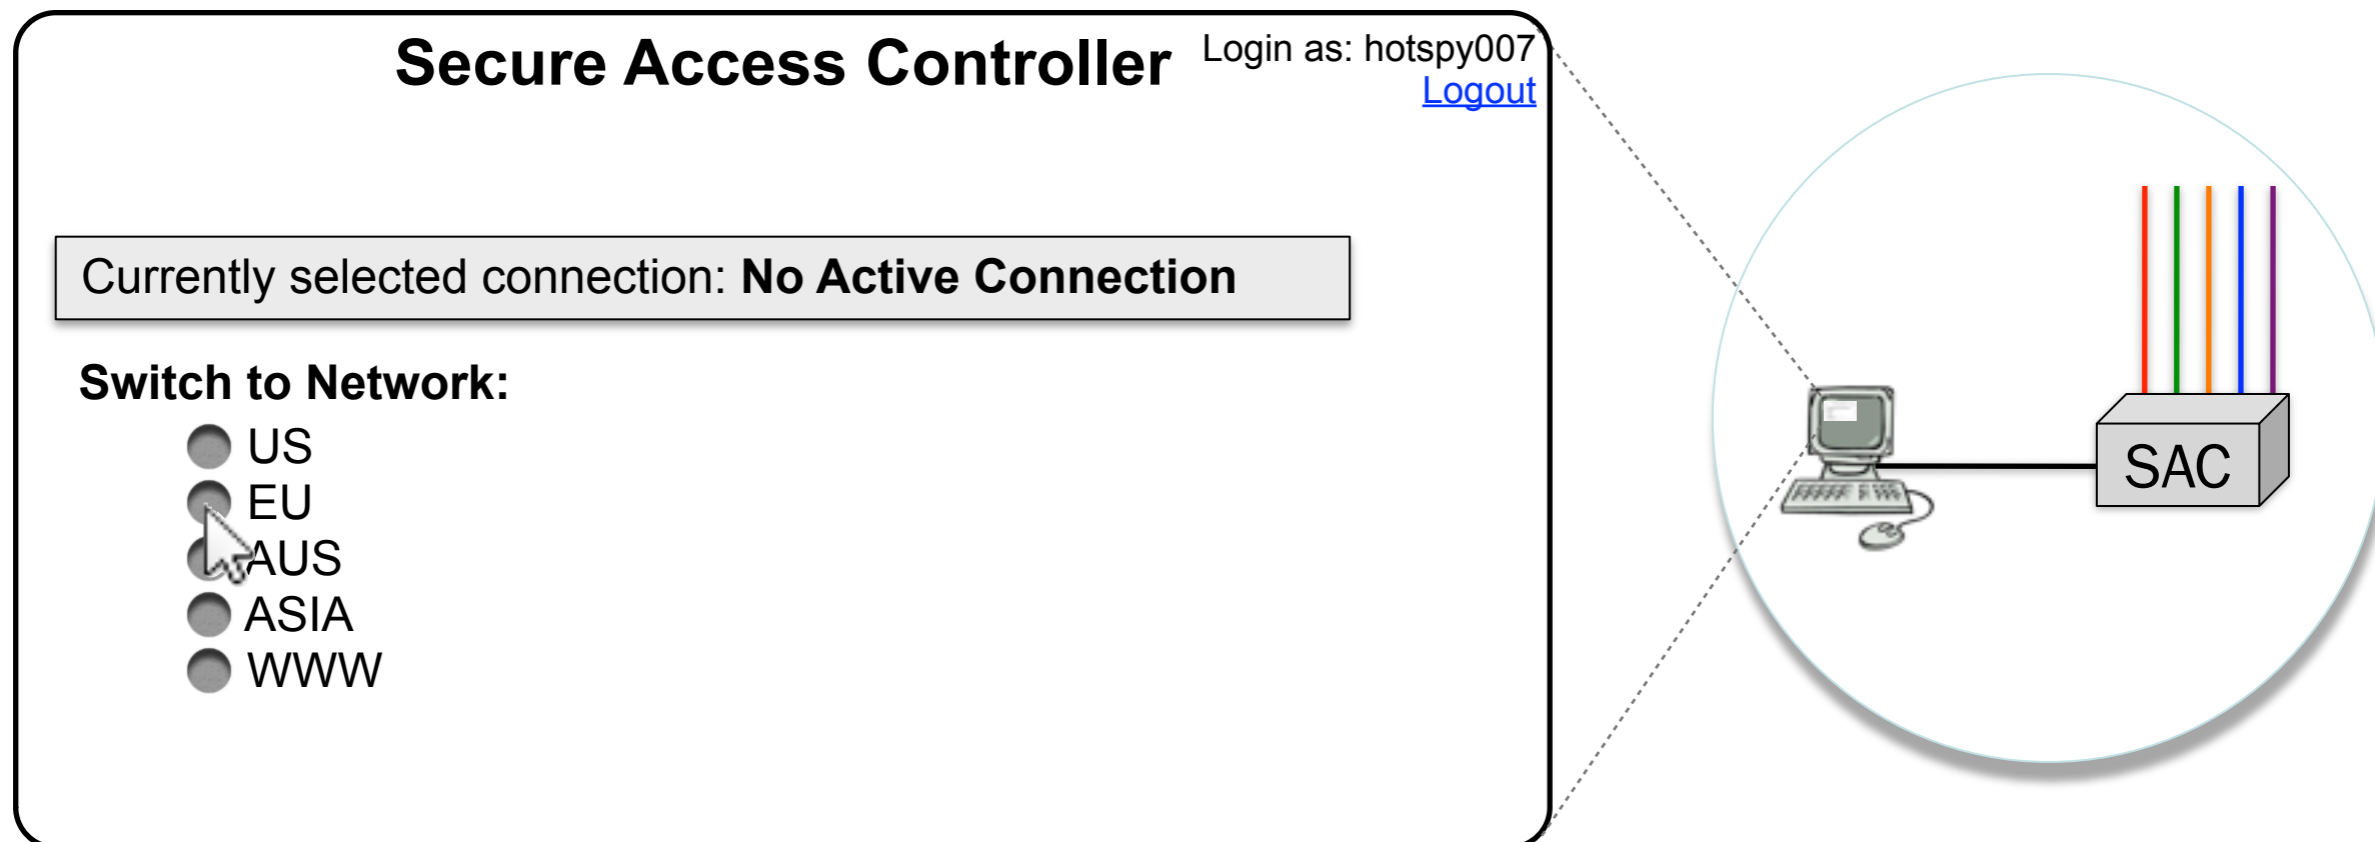
# Example System

- **Multilevel Secure Access Device**



# Example System

- **Multilevel Secure Access Device**



# Example System


- Multilevel Secure Access Device

**Secure Access Controller** Login as: hotspy007 [Logout](#)

Currently selected connection: **No Active Connection**

**Switch to Network:**

- US
- EU
- AUS
- ASIA
- WWW



# Example System

- **Multilevel Secure Access Device**

### Secure Access Controller

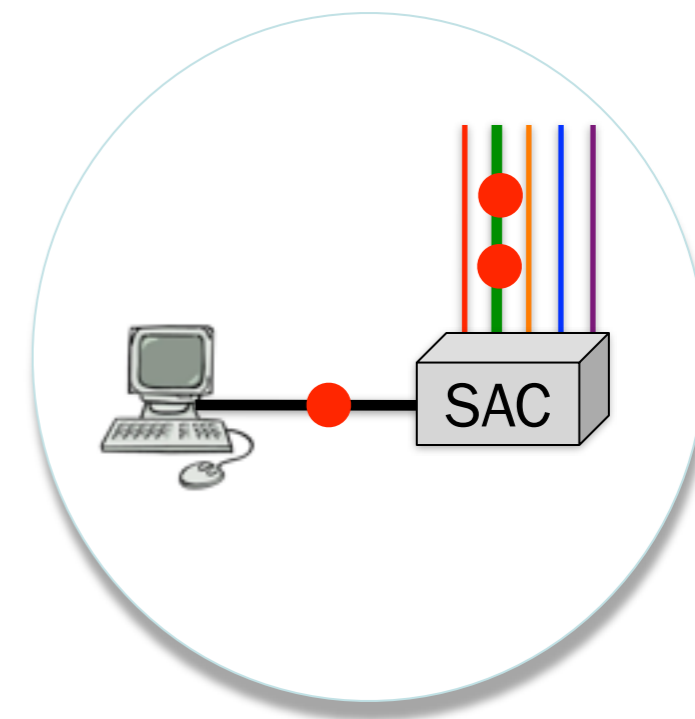
Login as: hotspy007 [Logout](#)

The Network has been successfully switched to **EU**

Currently selected connection: **EU**

**Switch to Network:**

- US
- EU
- AUS
- ASIA
- WWW



# Example System

- **Multilevel Secure Access Device**


**Secure Access Controller** Login as: hotspy007 [Logout](#)

The Network has been successfully switched to **EU**

Currently selected connection: **EU**

**Switch to Network:**

- US
- EU
- AUS
- ASIA
- WWW





# Example System

- **Multilevel Secure Access Device**


**Secure Access Controller** Login as: hotspy007 [Logout](#)

The Network has been successfully switched to **EU**

Currently selected connection: **EU**

**Switch to Network:**

- US
- EU
- AUS
- ASIA
- WWW



# Example System

- **Multilevel Secure Access Device**

### Secure Access Controller


Login as: hotspy007 [Logout](#)

The Network has been successfully switched to **WWW**

Currently selected connection: **WWW**

**Switch to Network:**


- US
- EU
- AUS
- ASIA
- WWW



# SAC System



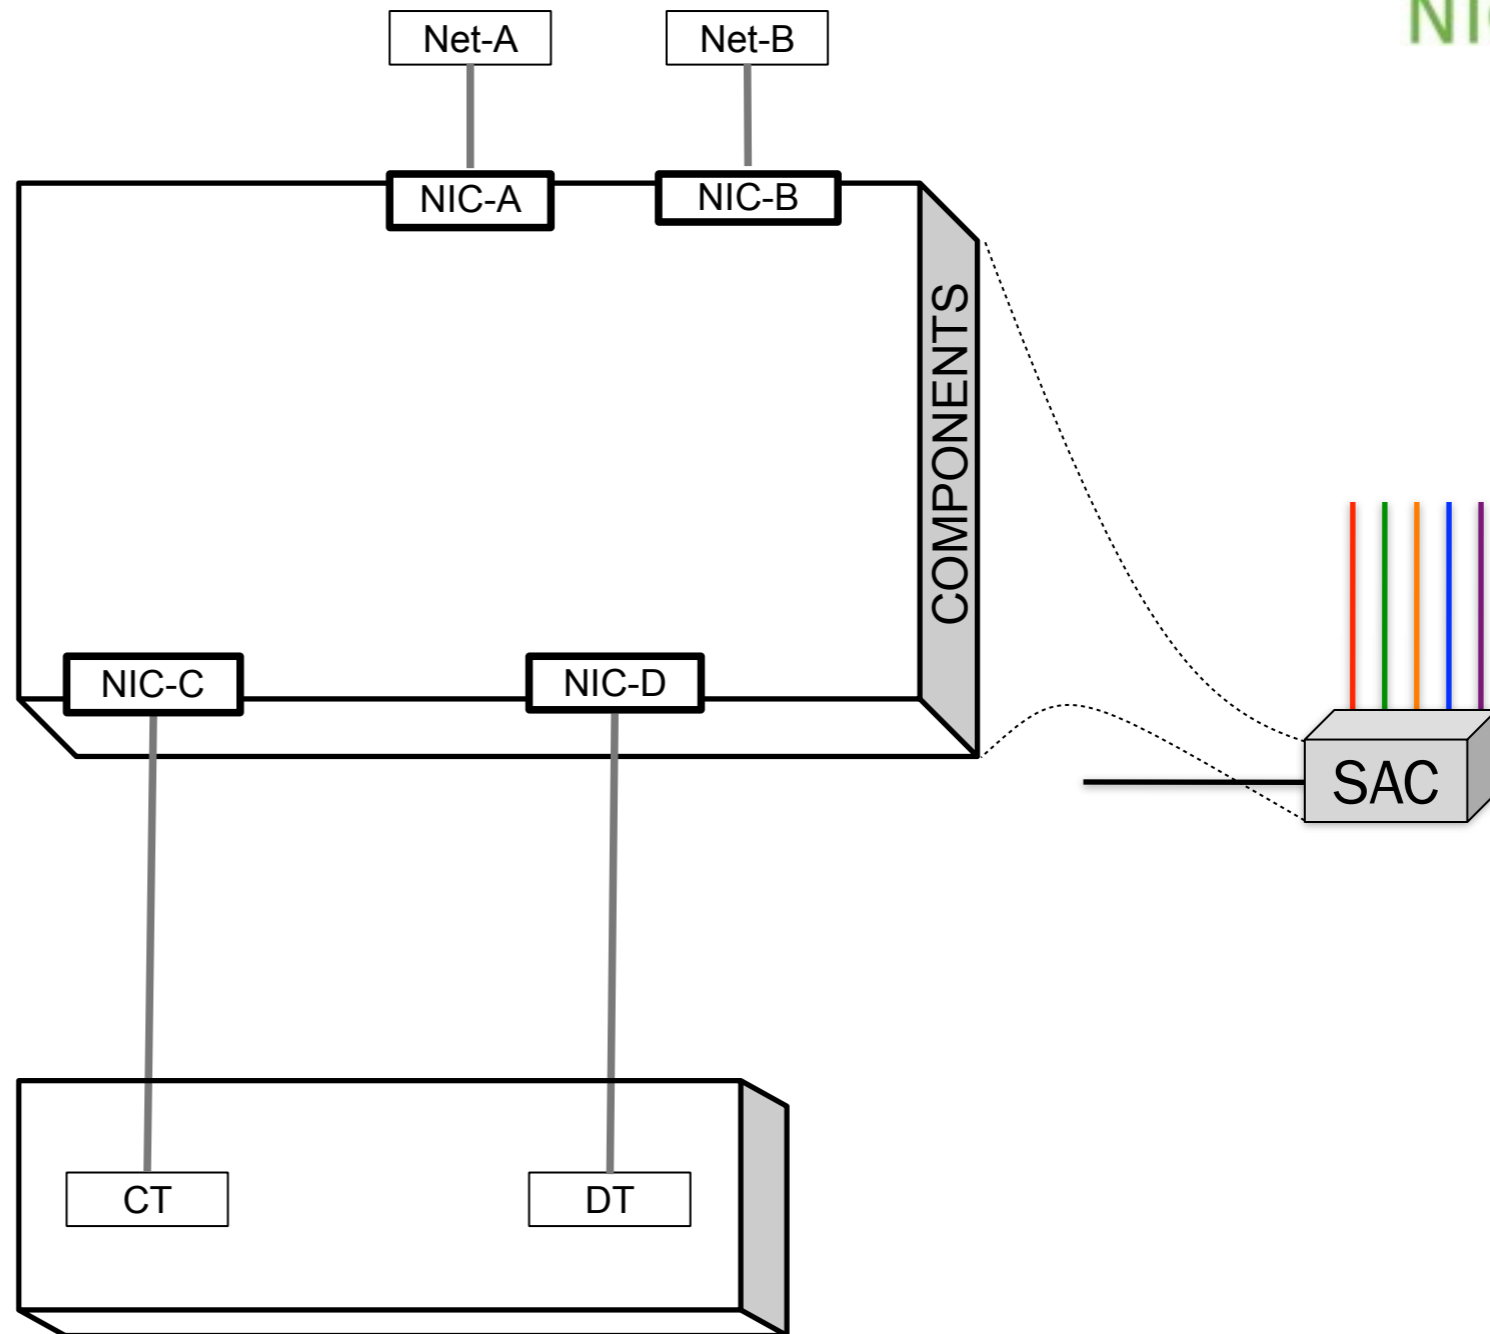
## Components



# SAC System



## Components



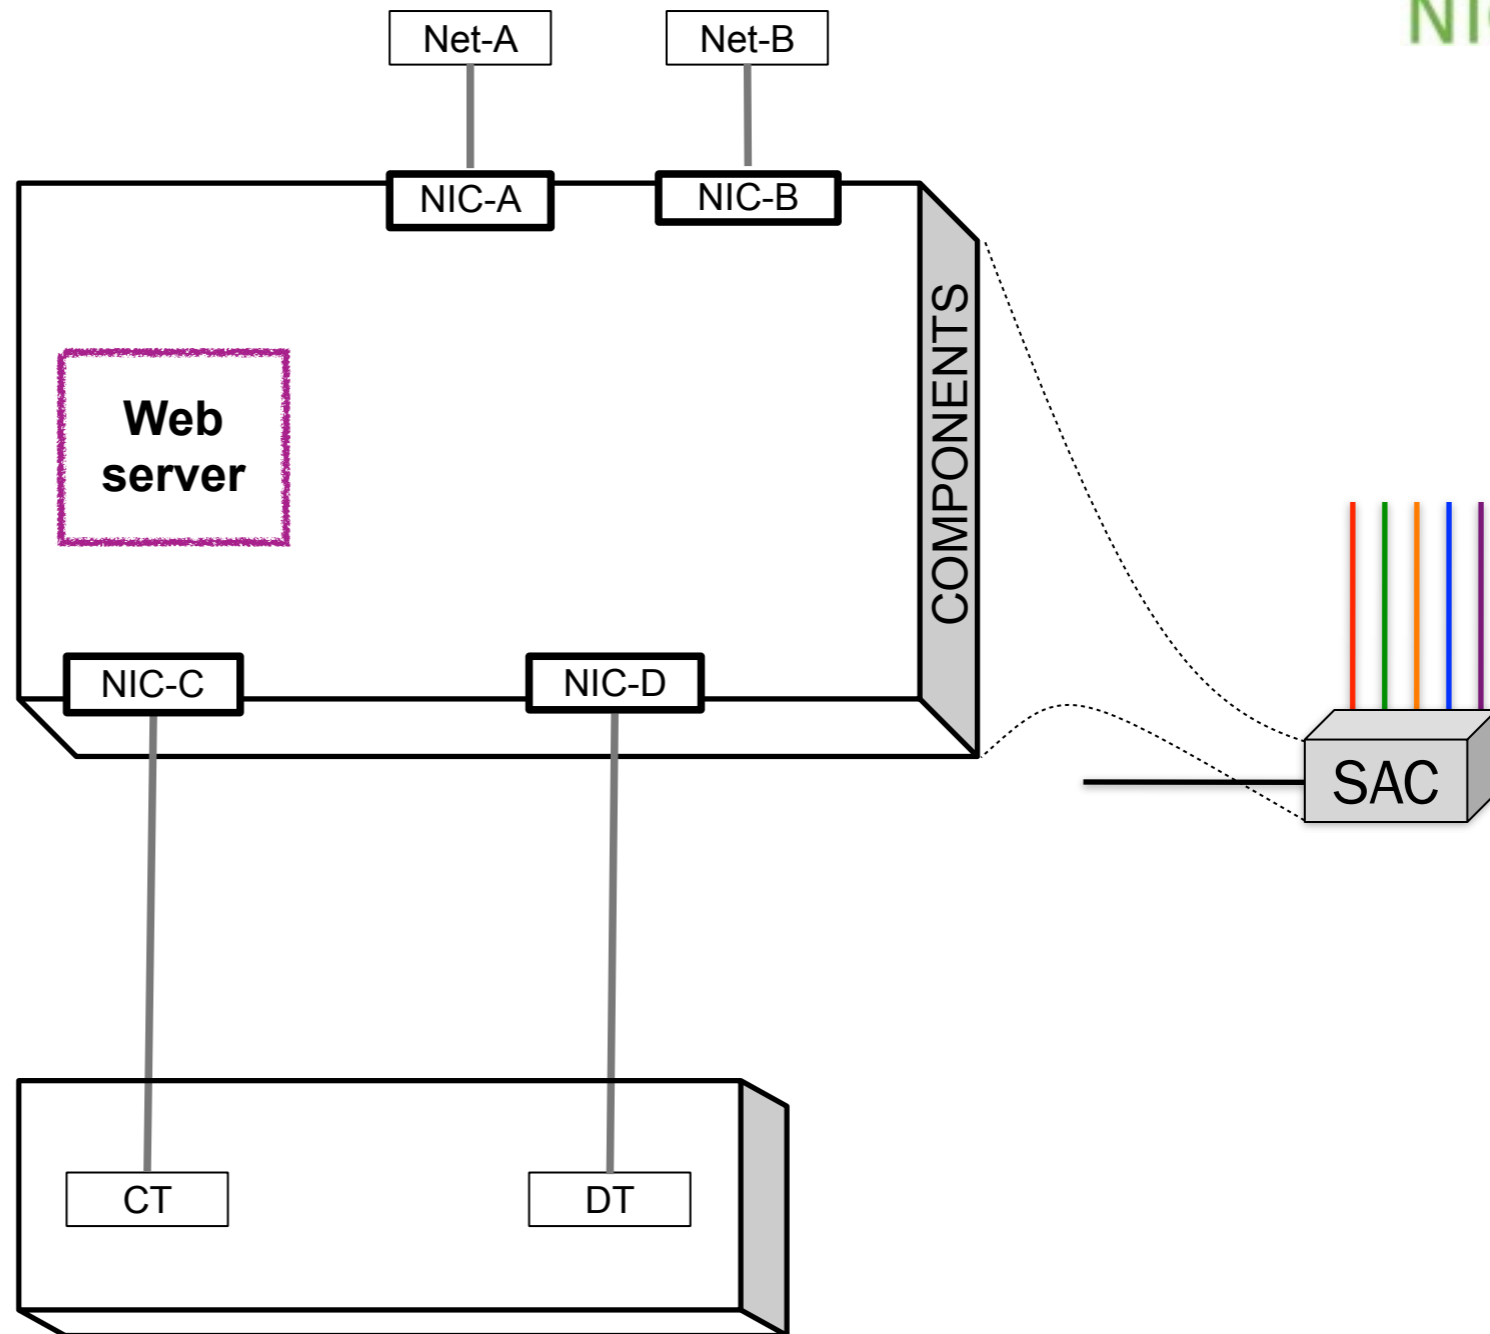
Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System



## Components



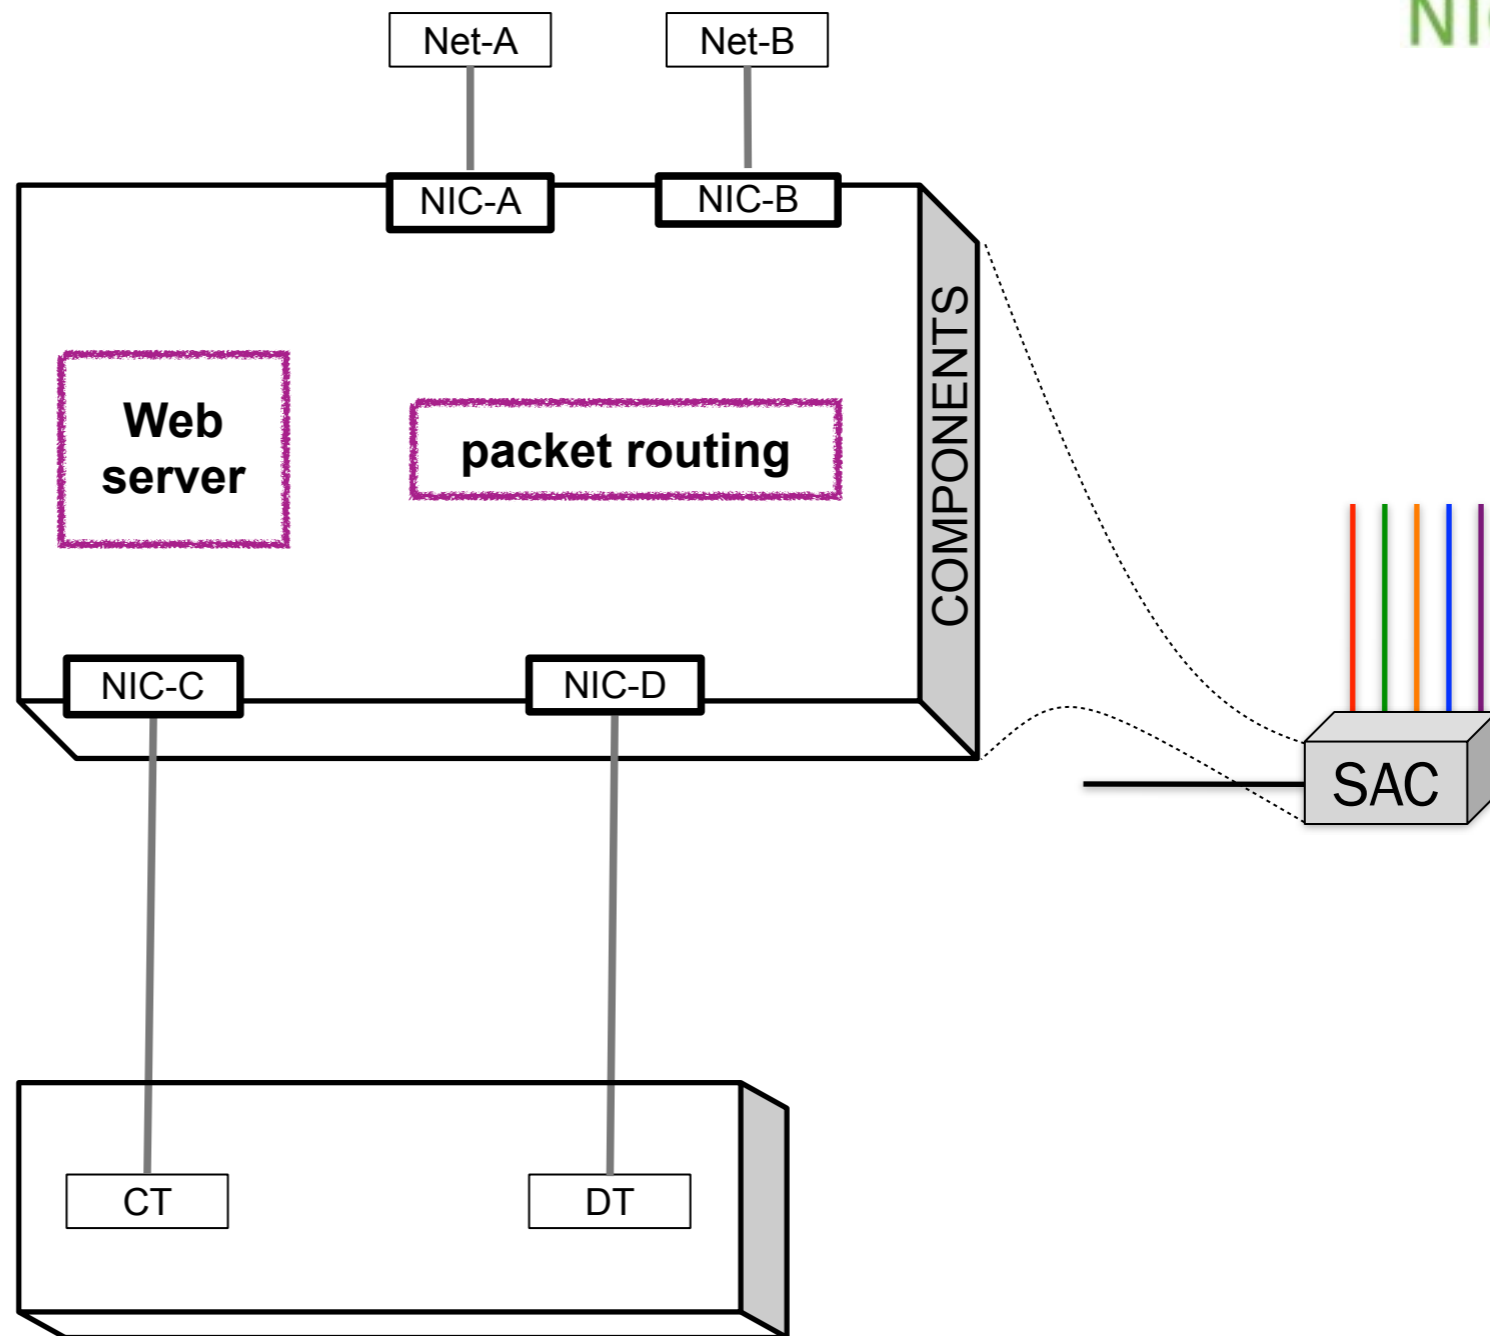
Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System



## Components



Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

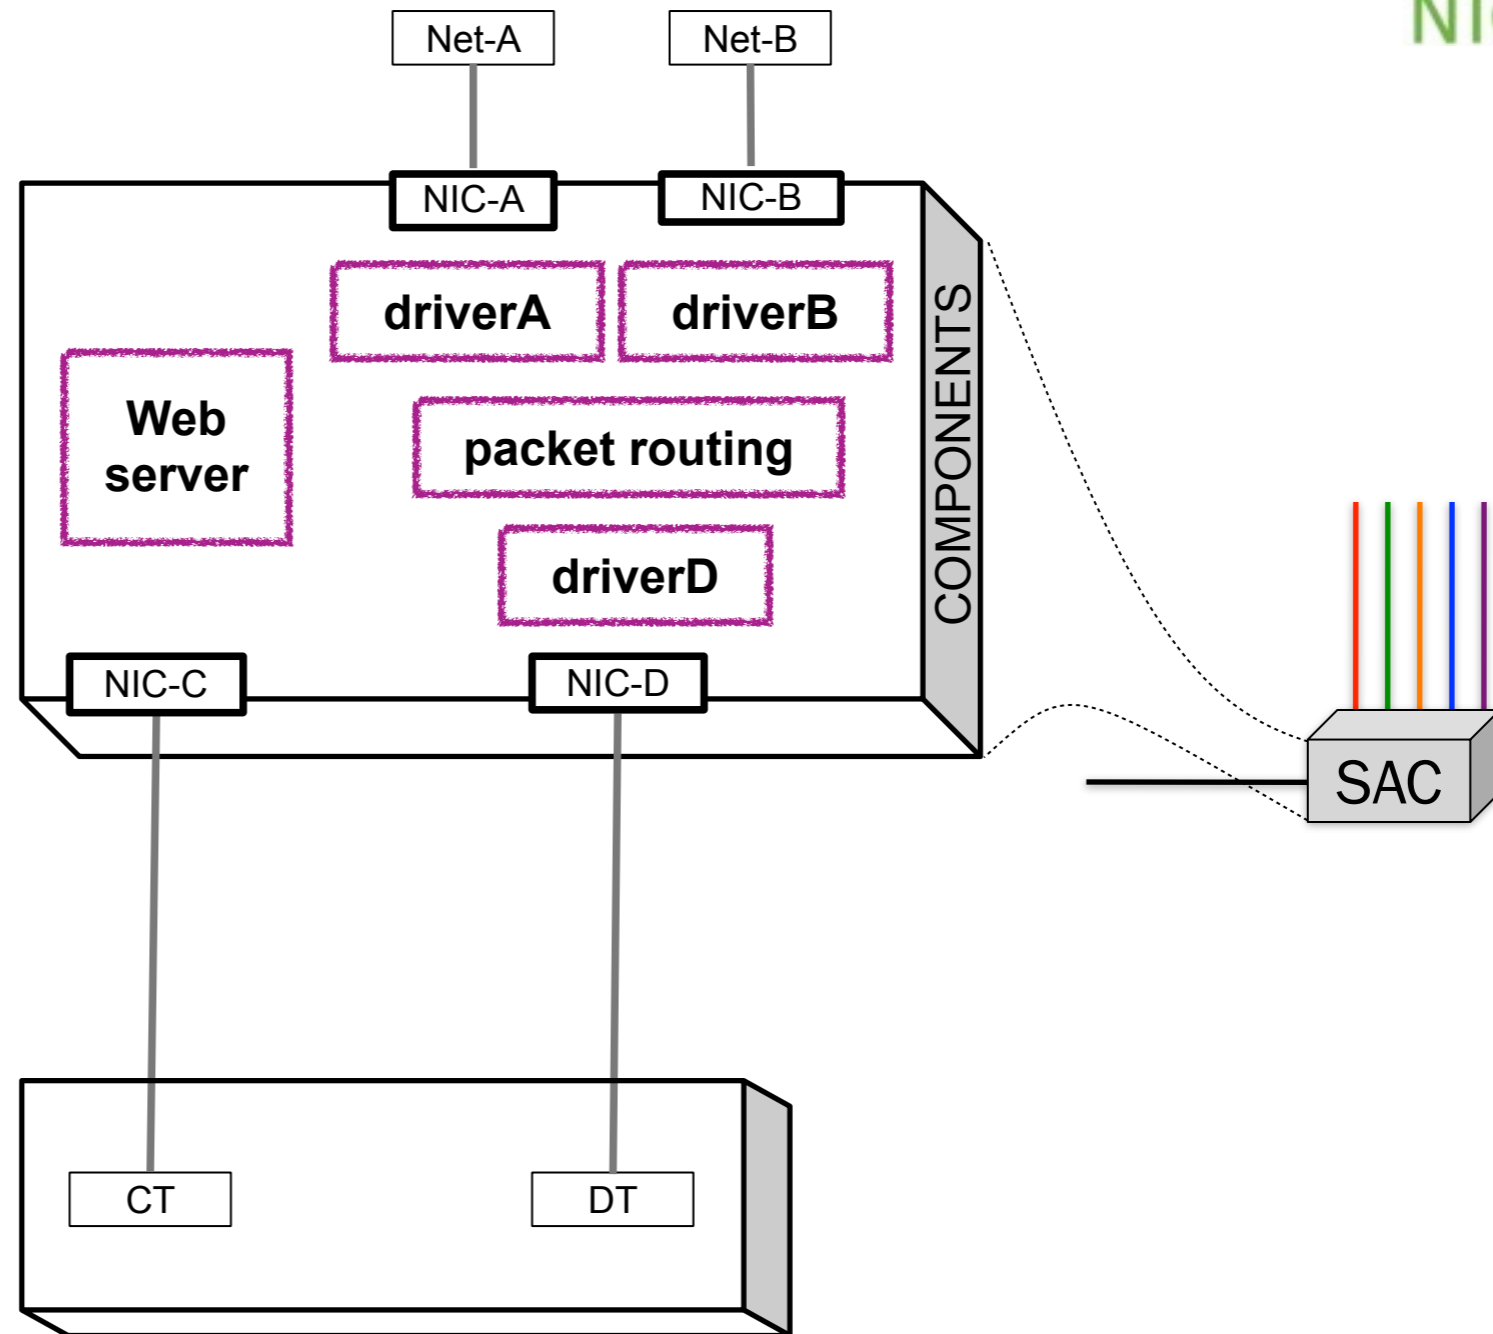
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal



# SAC System



## Components



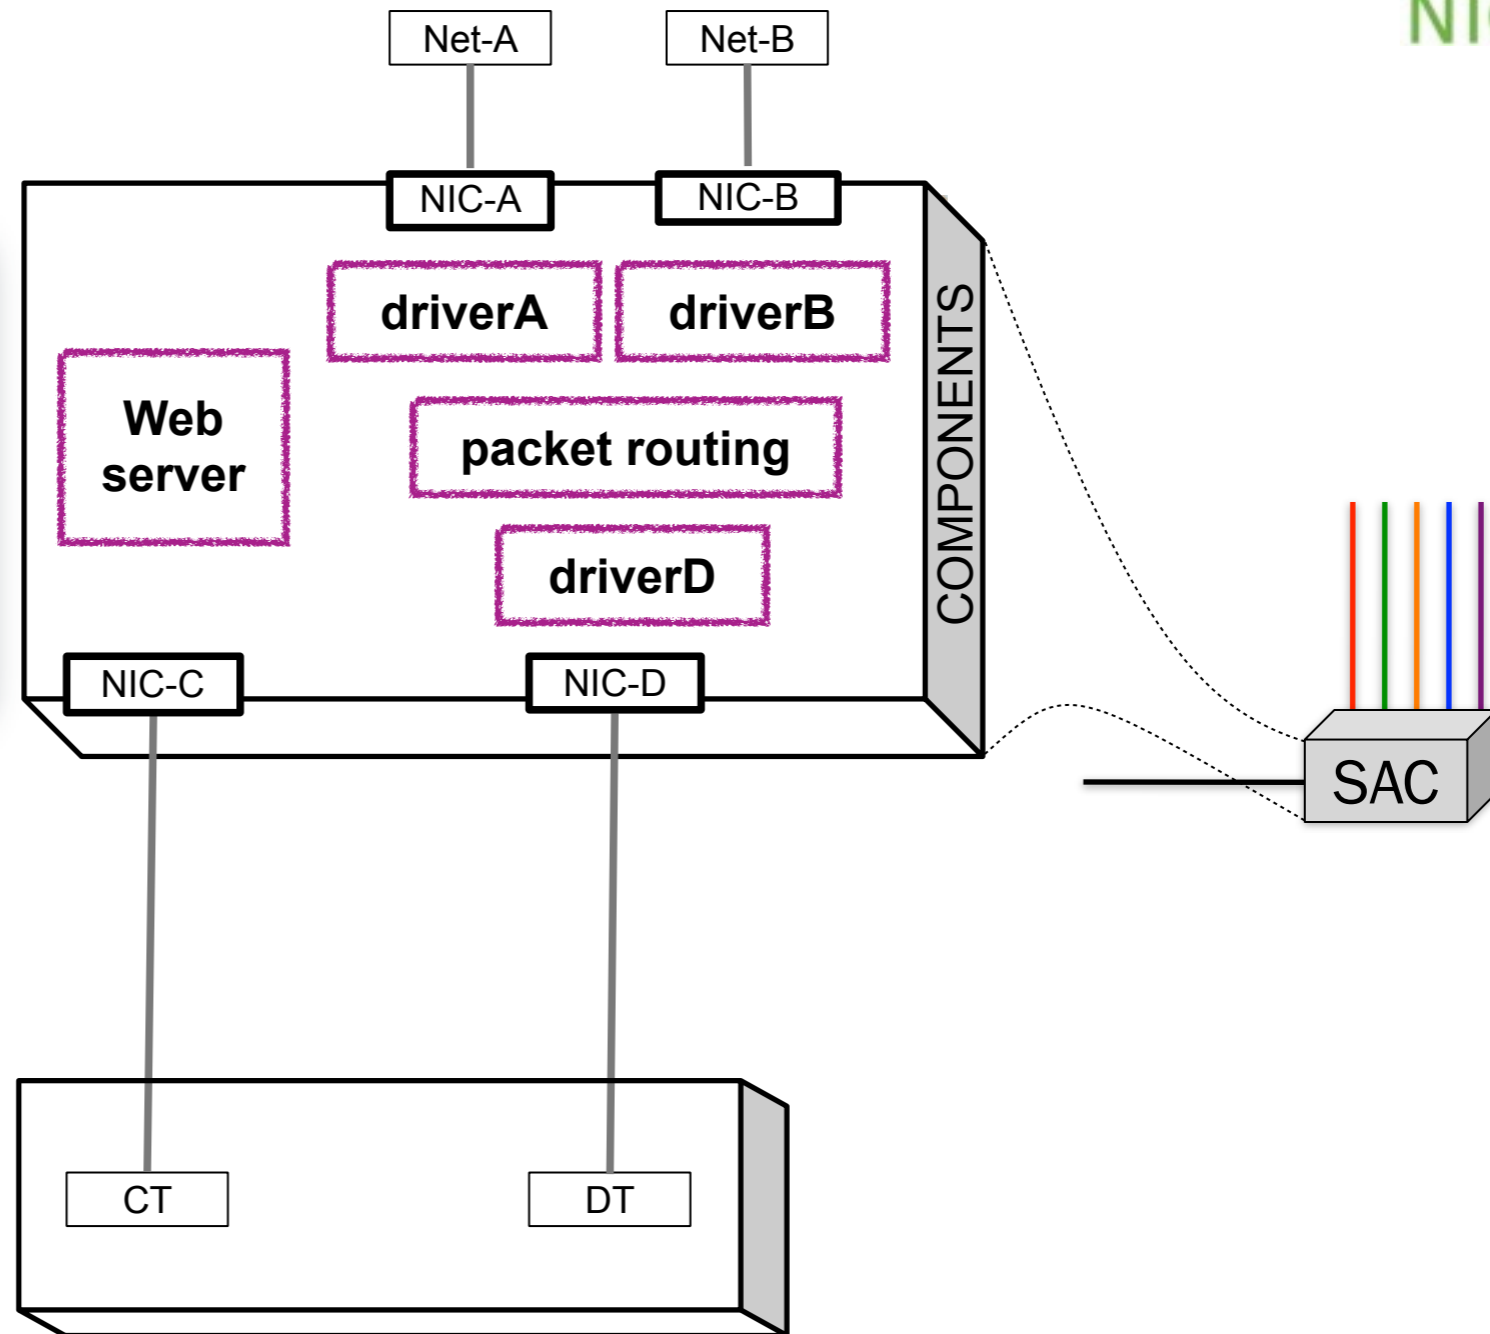
Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System

## Desired Property

No information flow between providers **A** and **B** through SAC even if they collaborate



Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

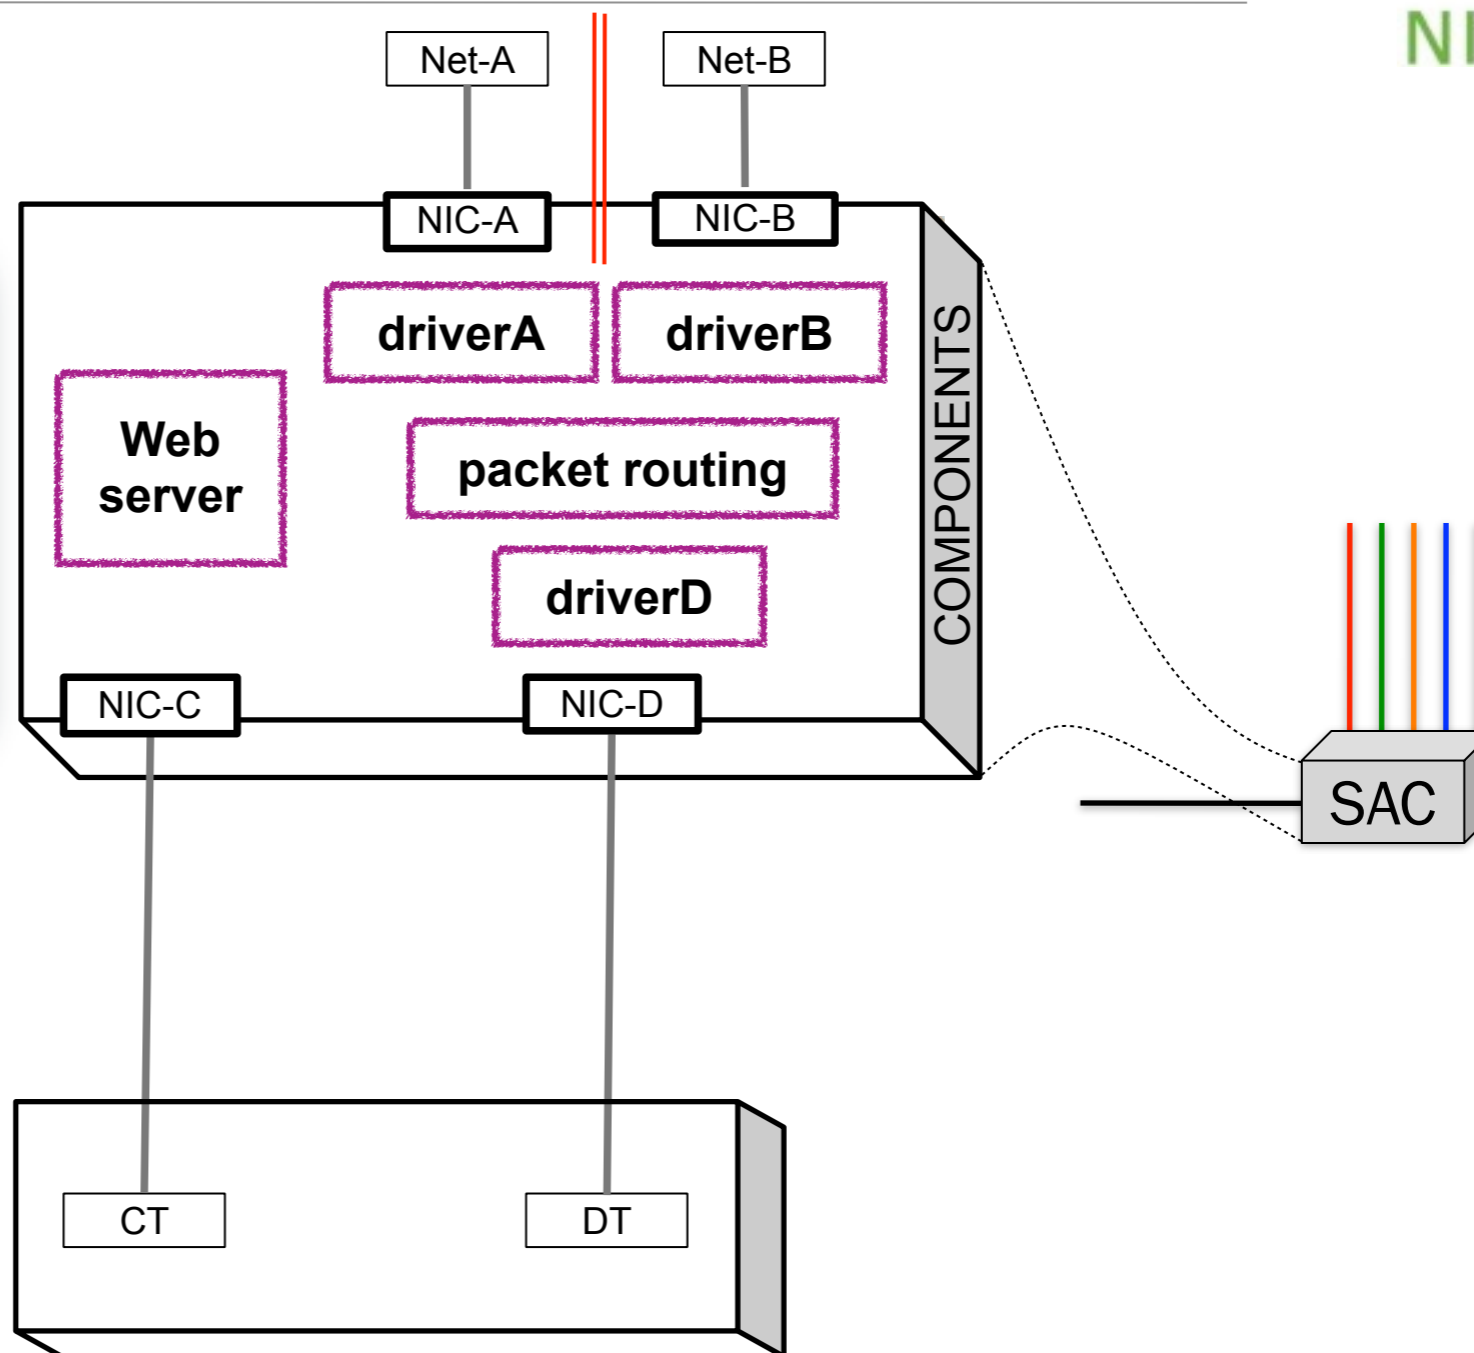
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System



## Desired Property

No information flow between providers **A** and **B** through SAC even if they collaborate



Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

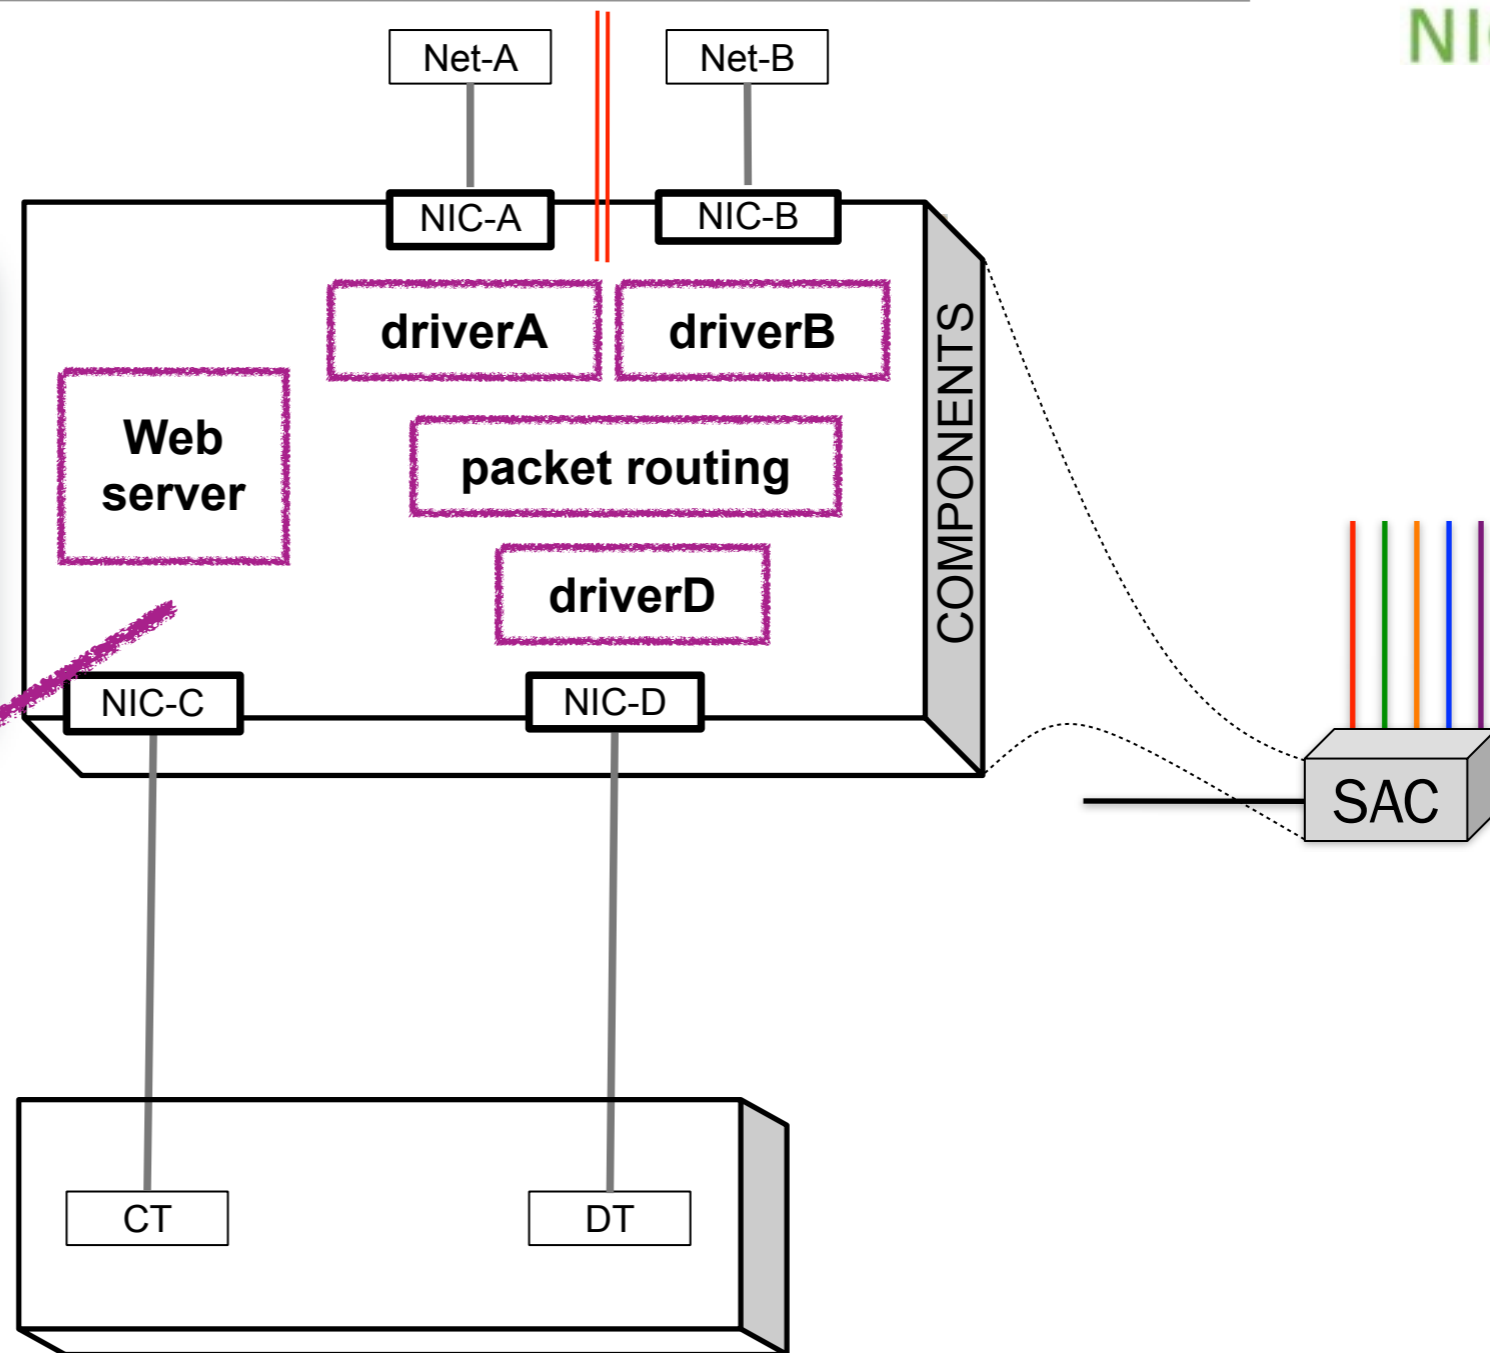
# SAC System



## Desired Property

No information flow between providers **A** and **B** through SAC even if they collaborate

Proving all this correct?



Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

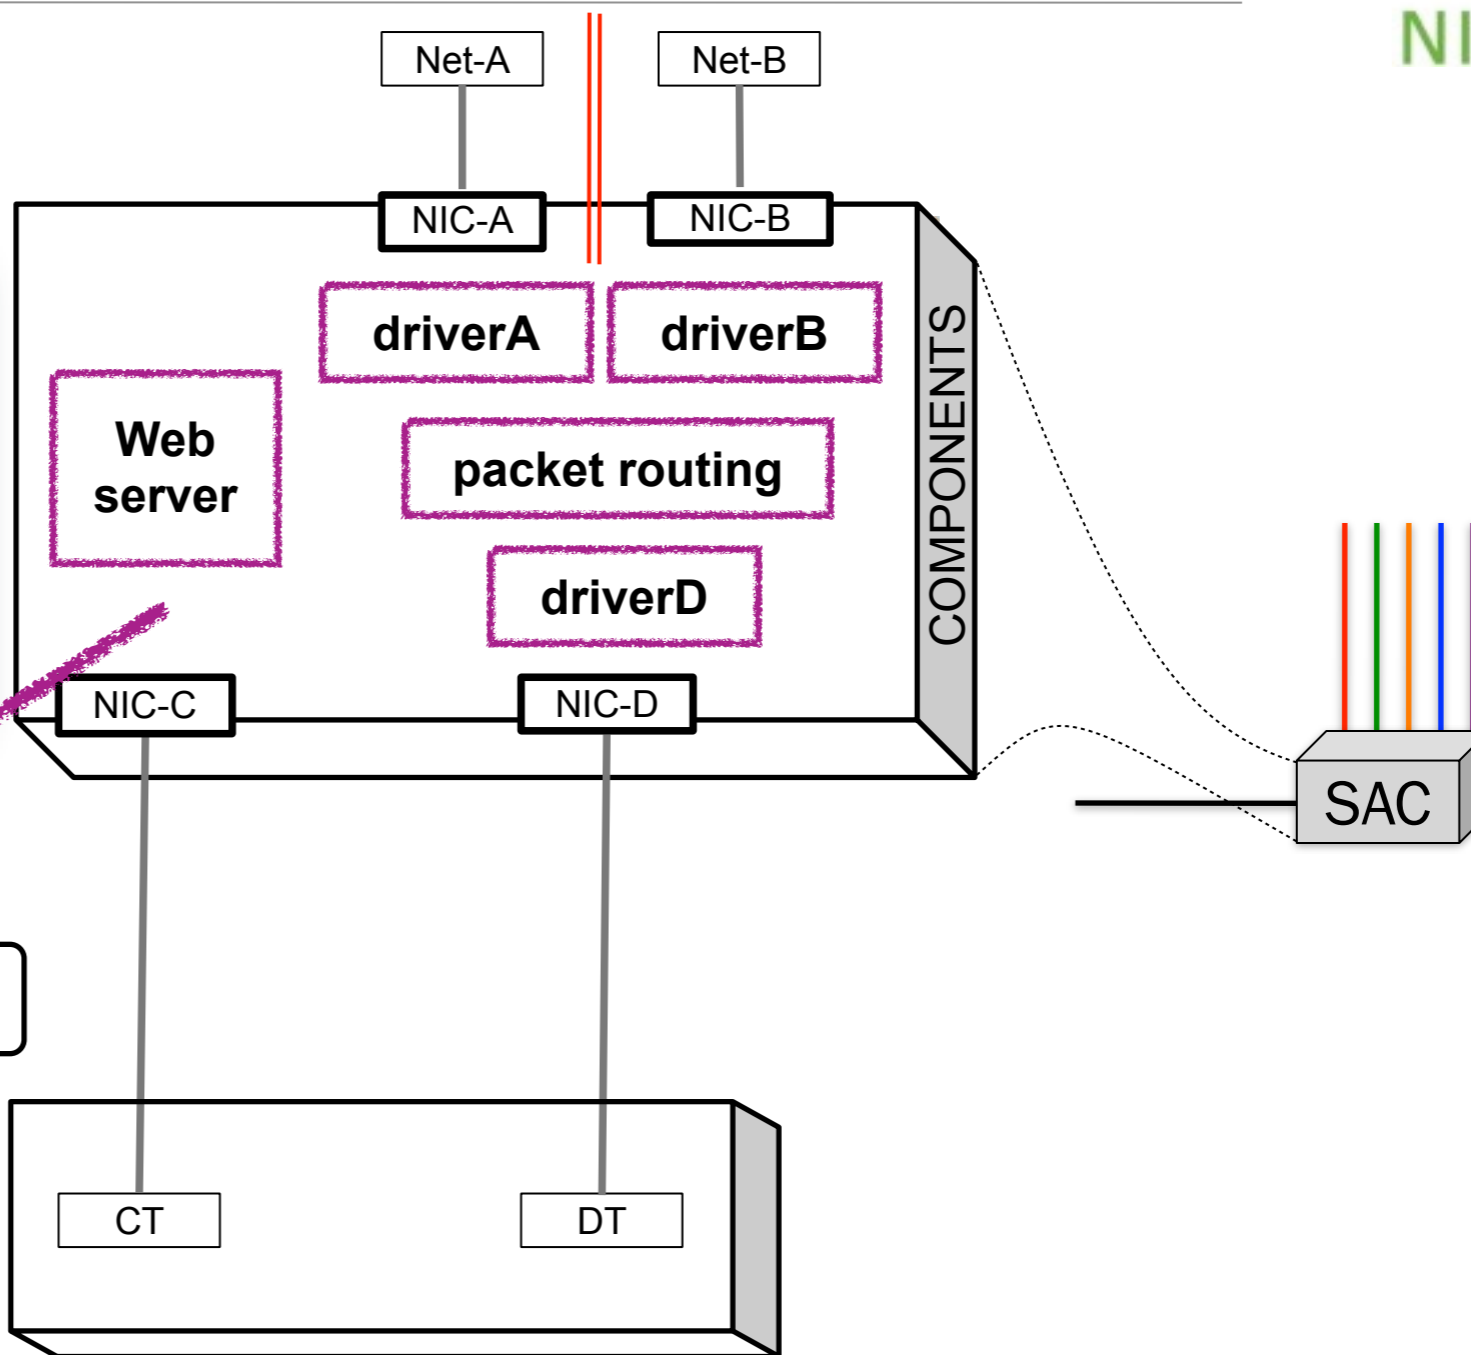
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System



## Desired Property

No information flow between providers **A** and **B** through SAC even if they collaborate



Proving all this correct?

**NO!**

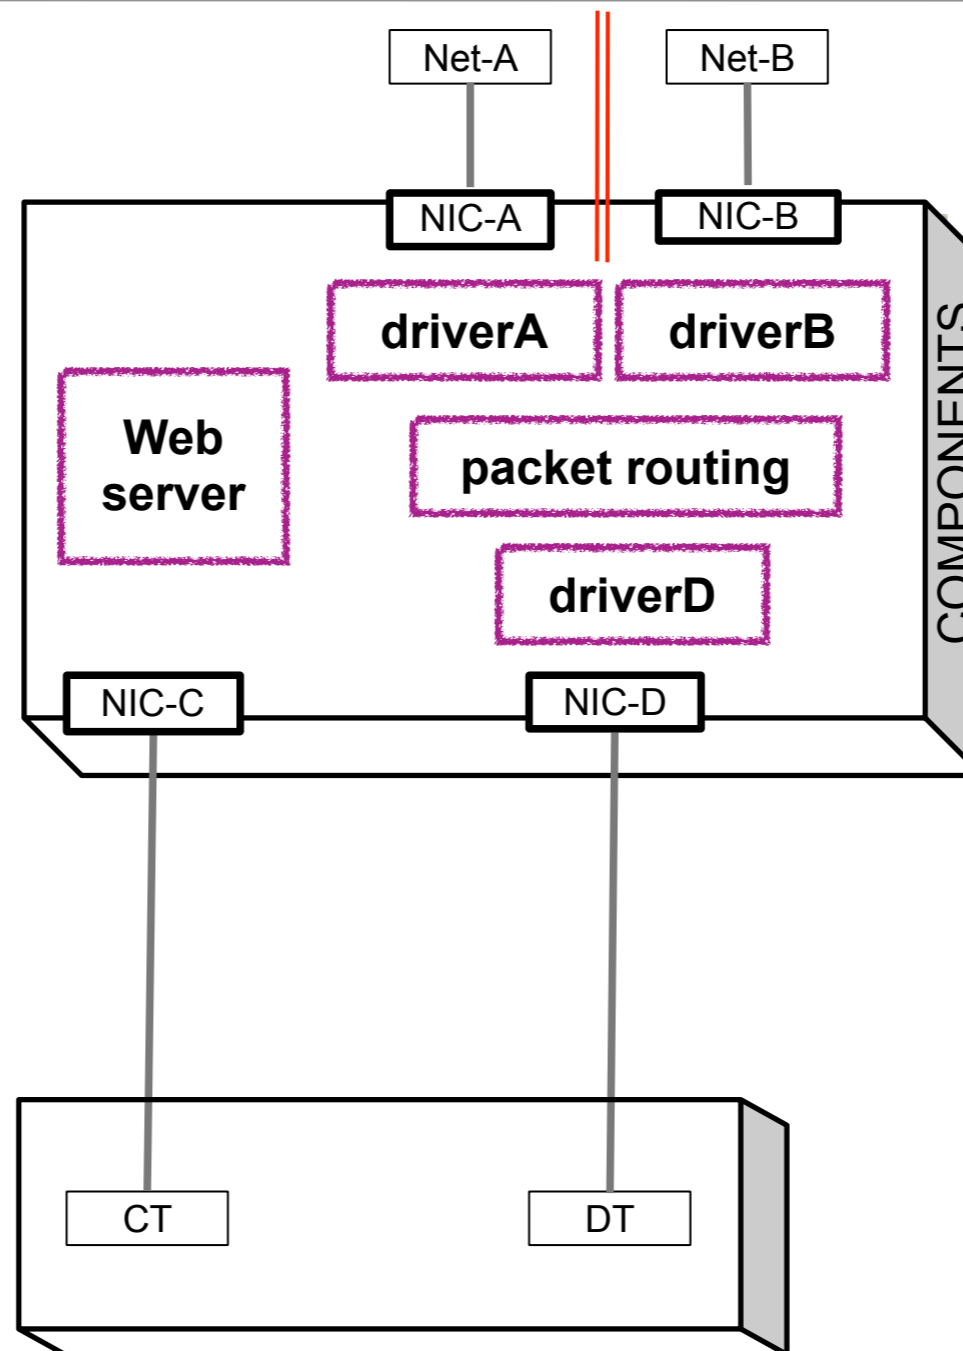


Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System

## Design



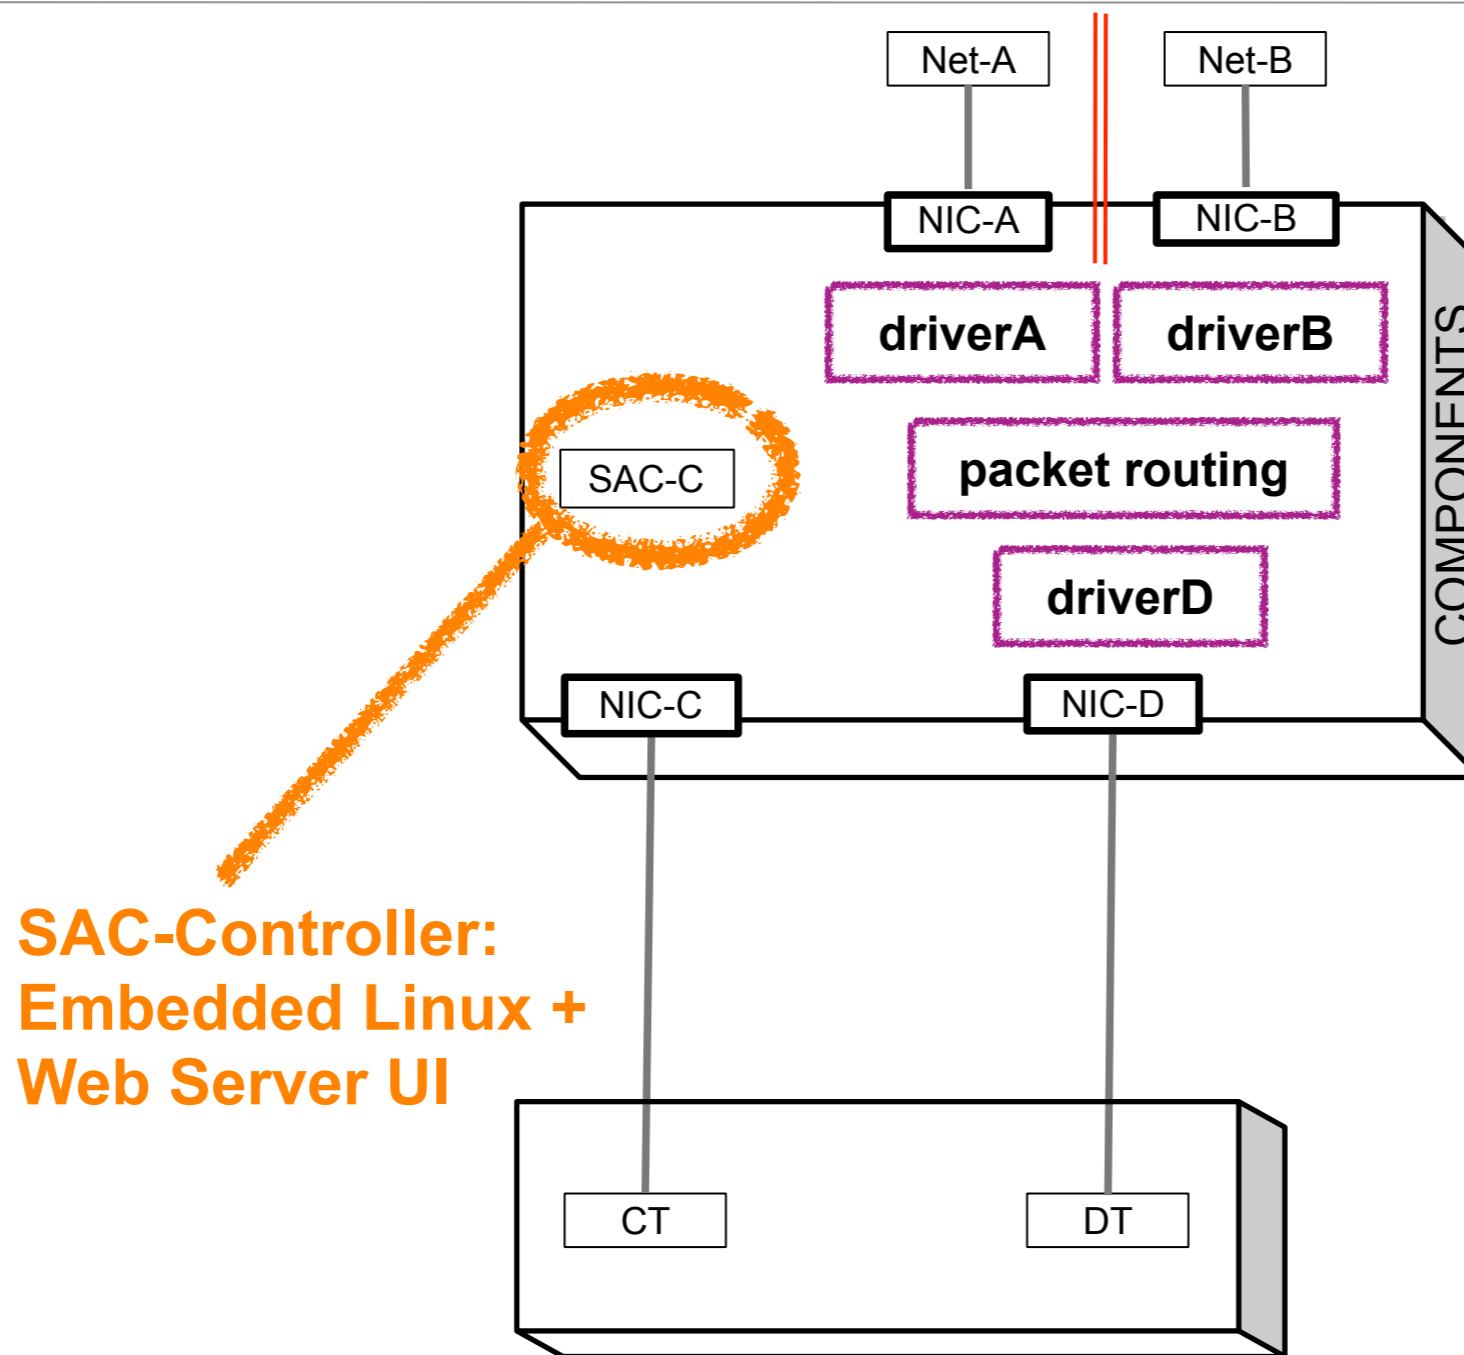
Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal



# SAC System

## Design



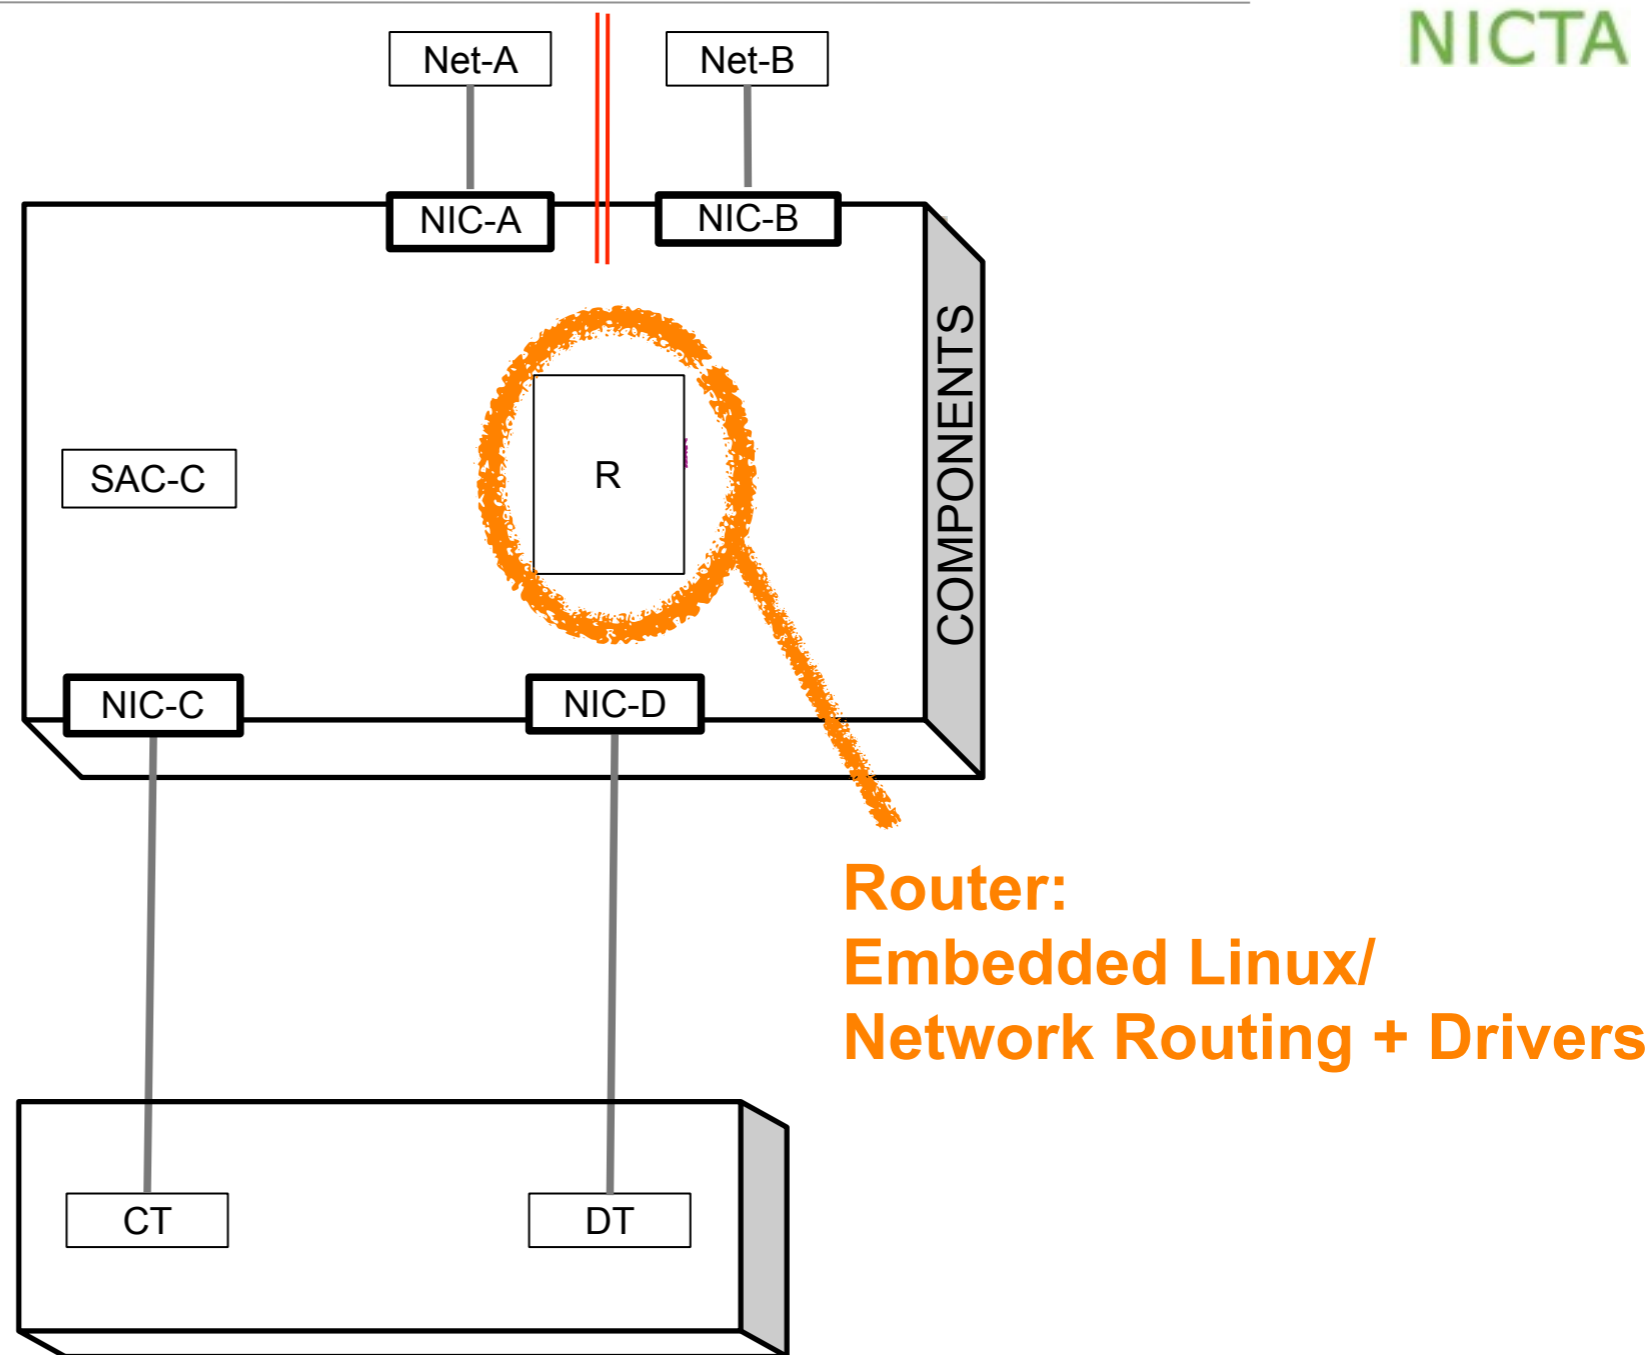
**SAC-Controller:  
Embedded Linux +  
Web Server UI**

Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

# SAC System

## Design



Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

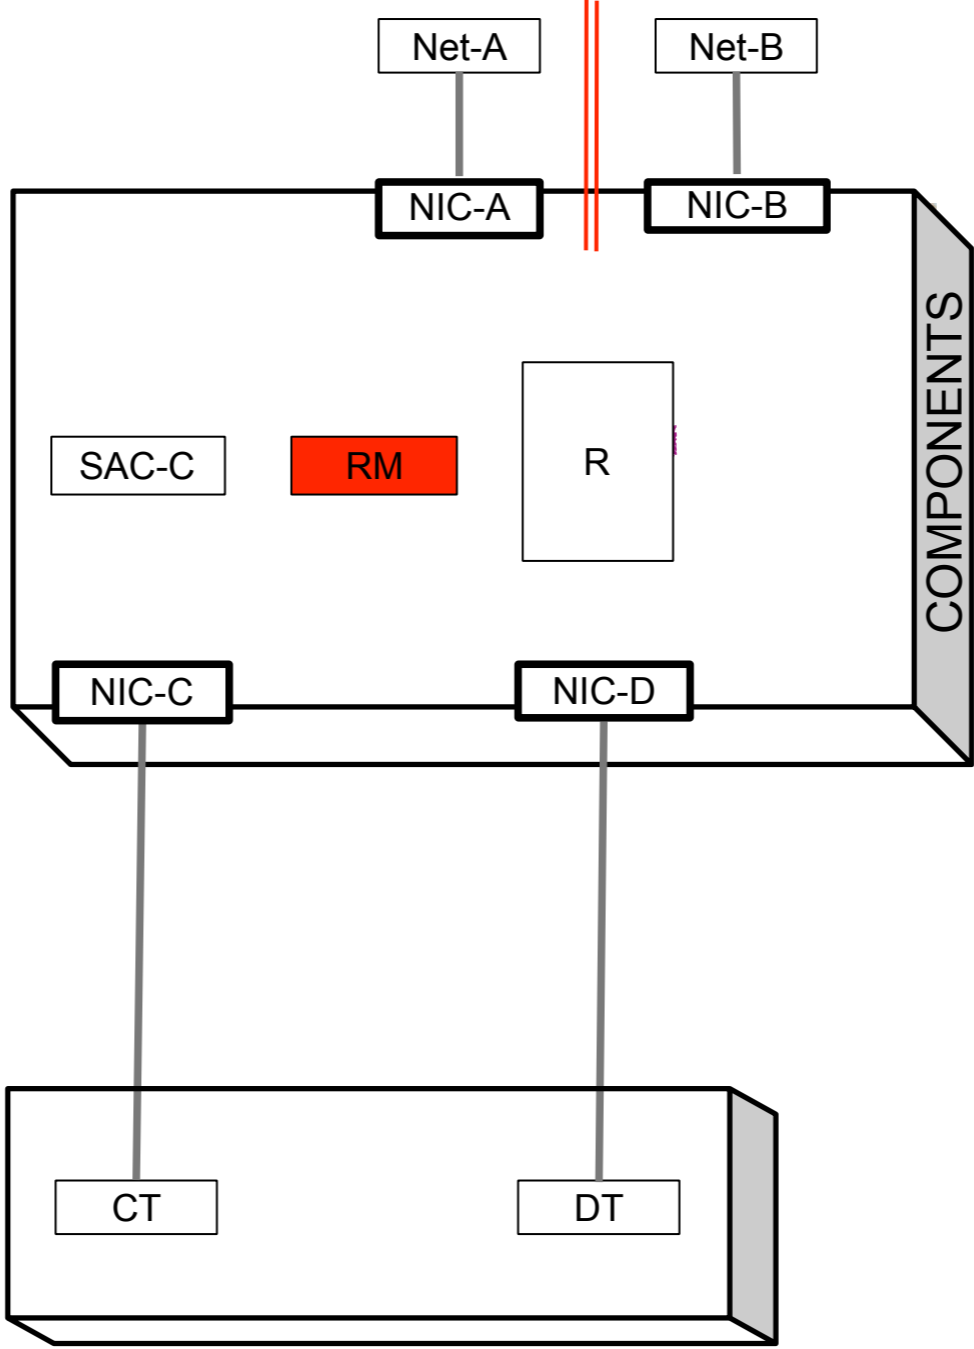
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

R = Router  
RM = Router Manager  
SAC-C = SAC Controller

# SAC System



## Design

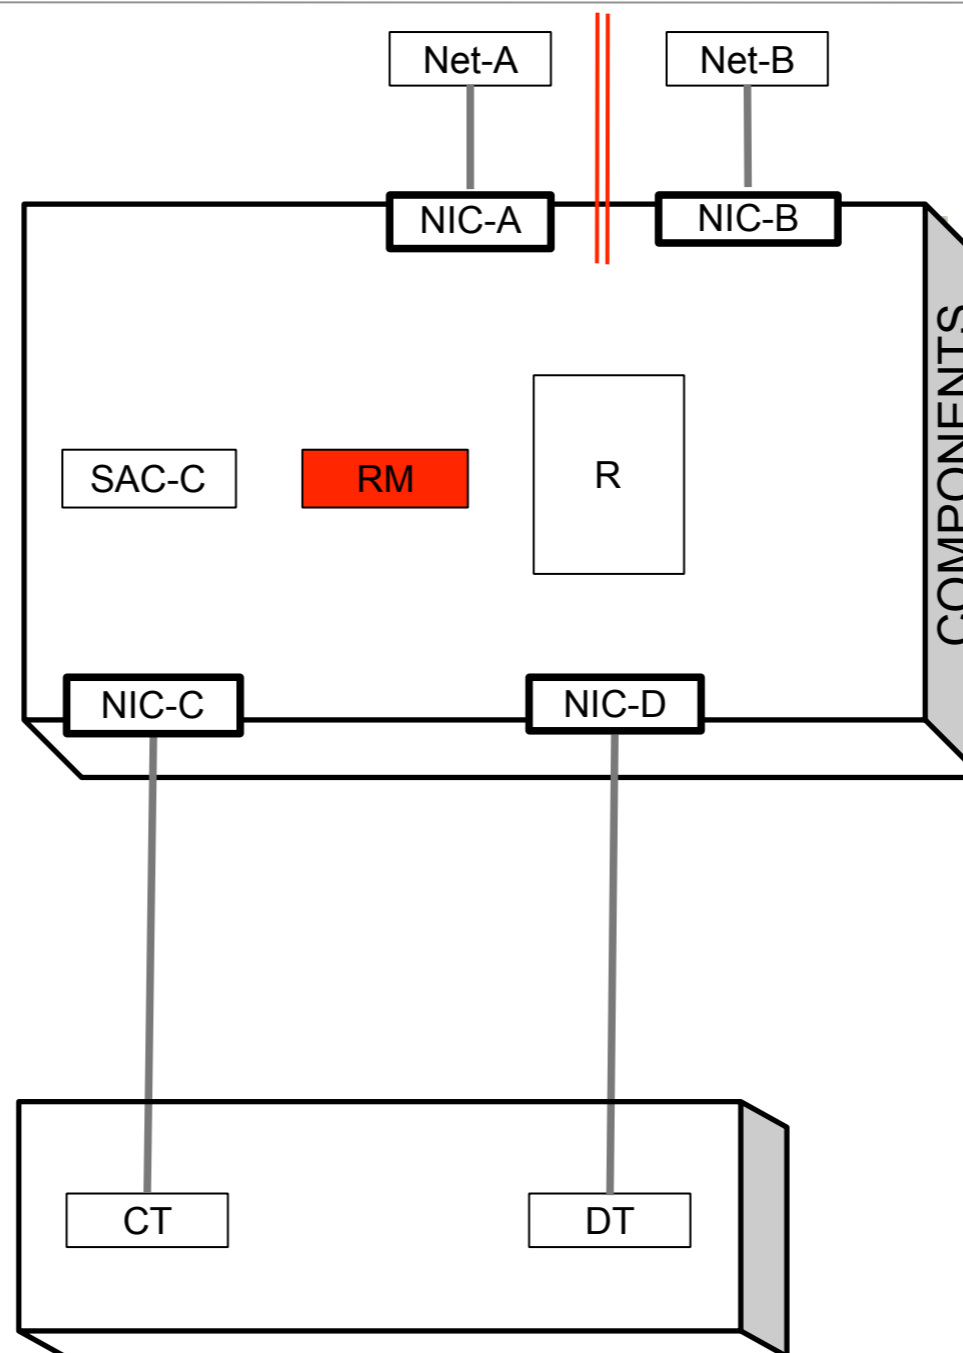


Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B  
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal  
R = Router  
RM = Router Manager  
SAC-C = SAC Controller

# SAC System



## Design



Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

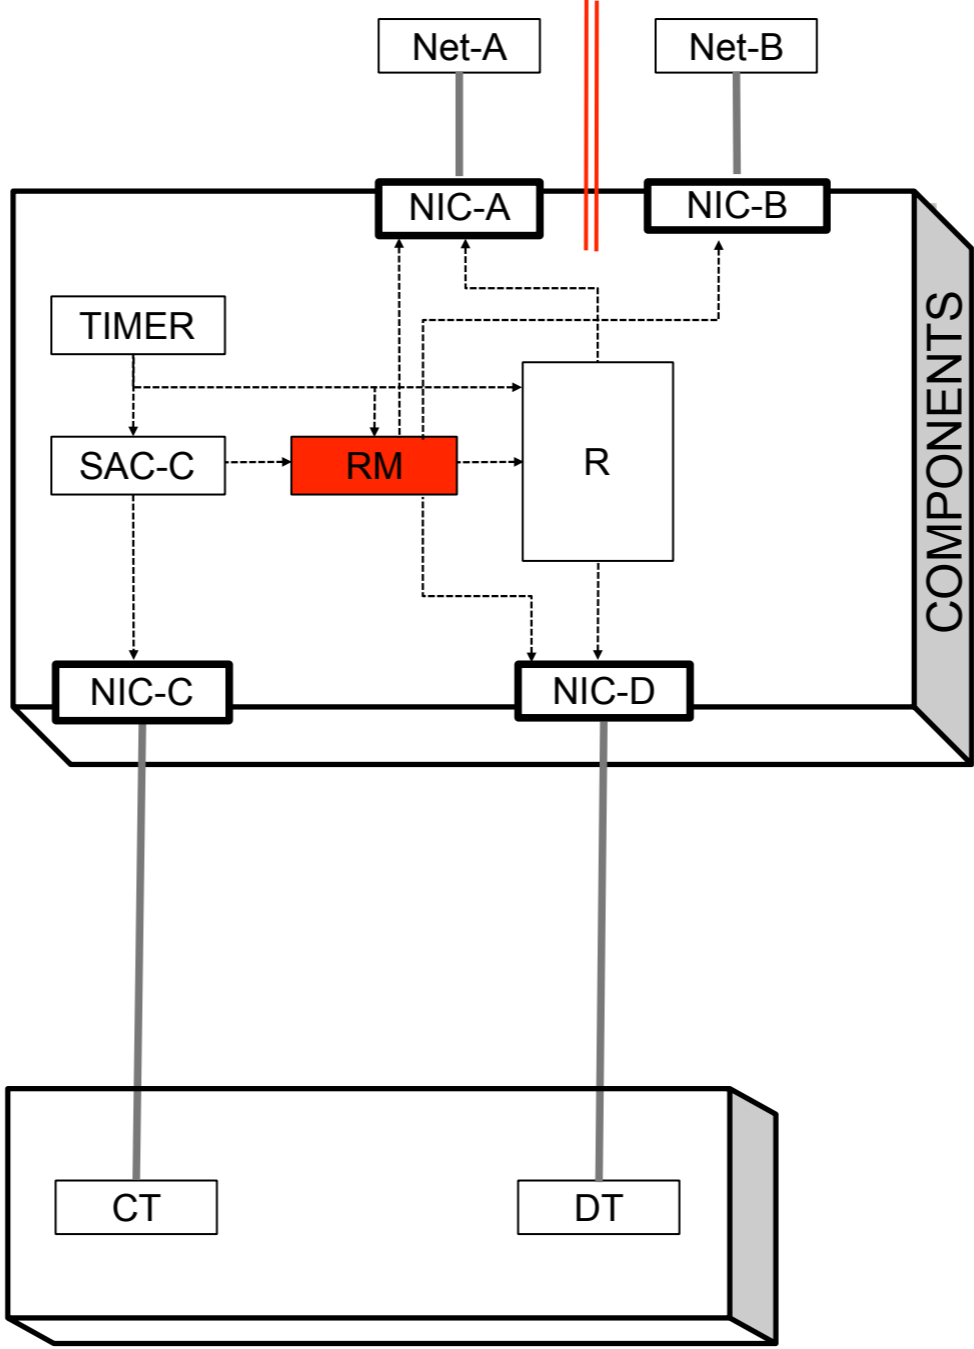
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

R = Router  
RM = Router Manager  
SAC-C = SAC Controller

# SAC System



## Design



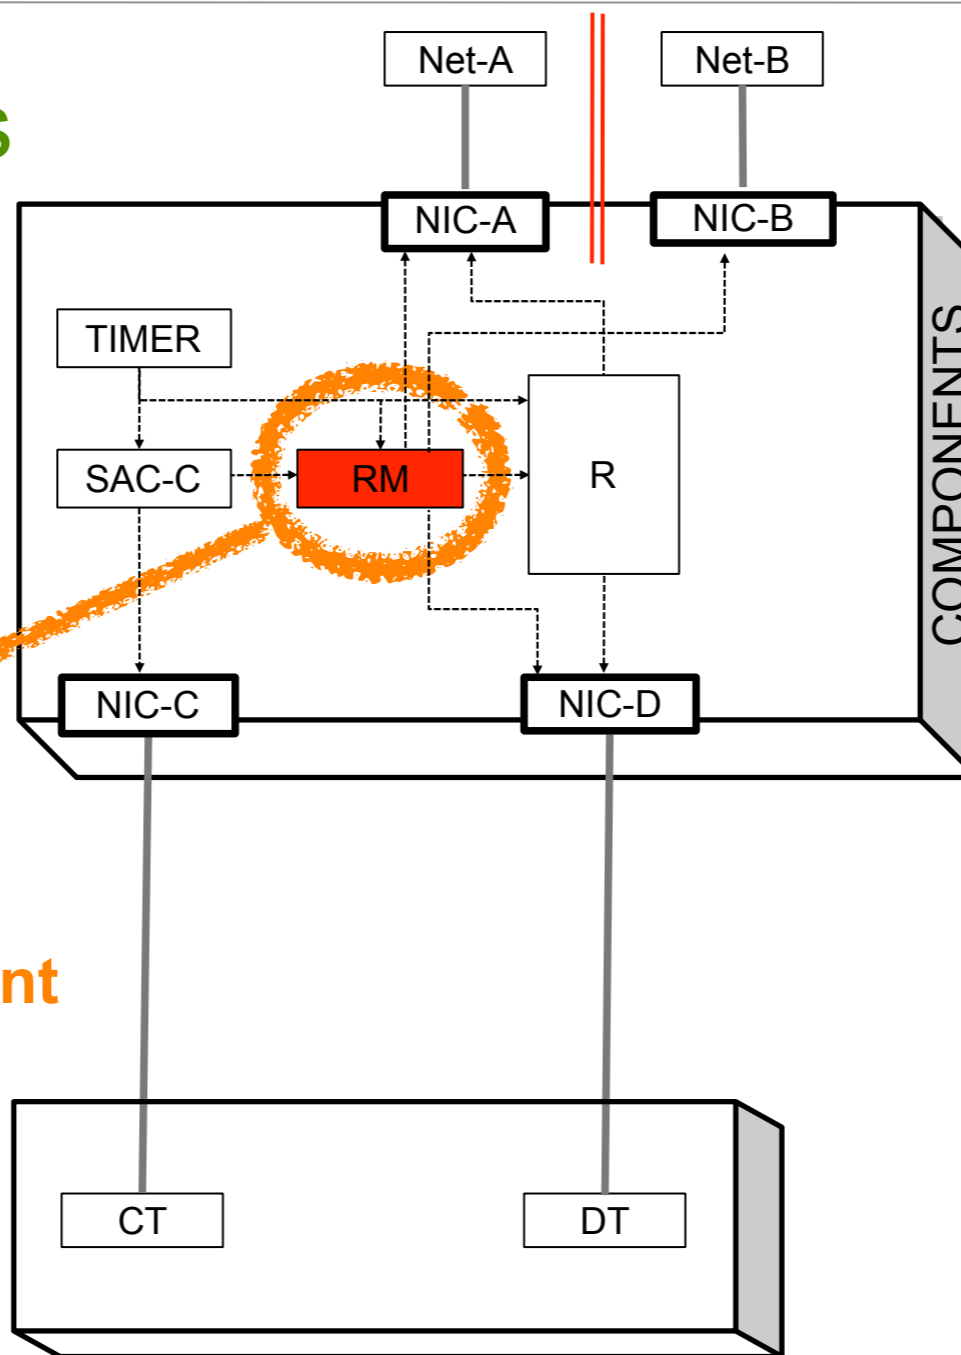
----->  
has access to

Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B  
NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal  
R= Router  
RM = Router Manager  
SAC-C = SAC Controller

# SAC System

## Trusted Components

**Router Manager:  
< 2kloc  
only trusted component**




Net-A = Network A  
Net-B = Network B  
NIC-A = Network Card for Network A  
NIC-B = Network Card for Network B

NIC-C = Control Network Card  
NIC-D = Data Network Card  
CT = Control Terminal  
DT = Data Terminal

R= Router  
RM = Router Manager  
SAC-C = SAC Controller

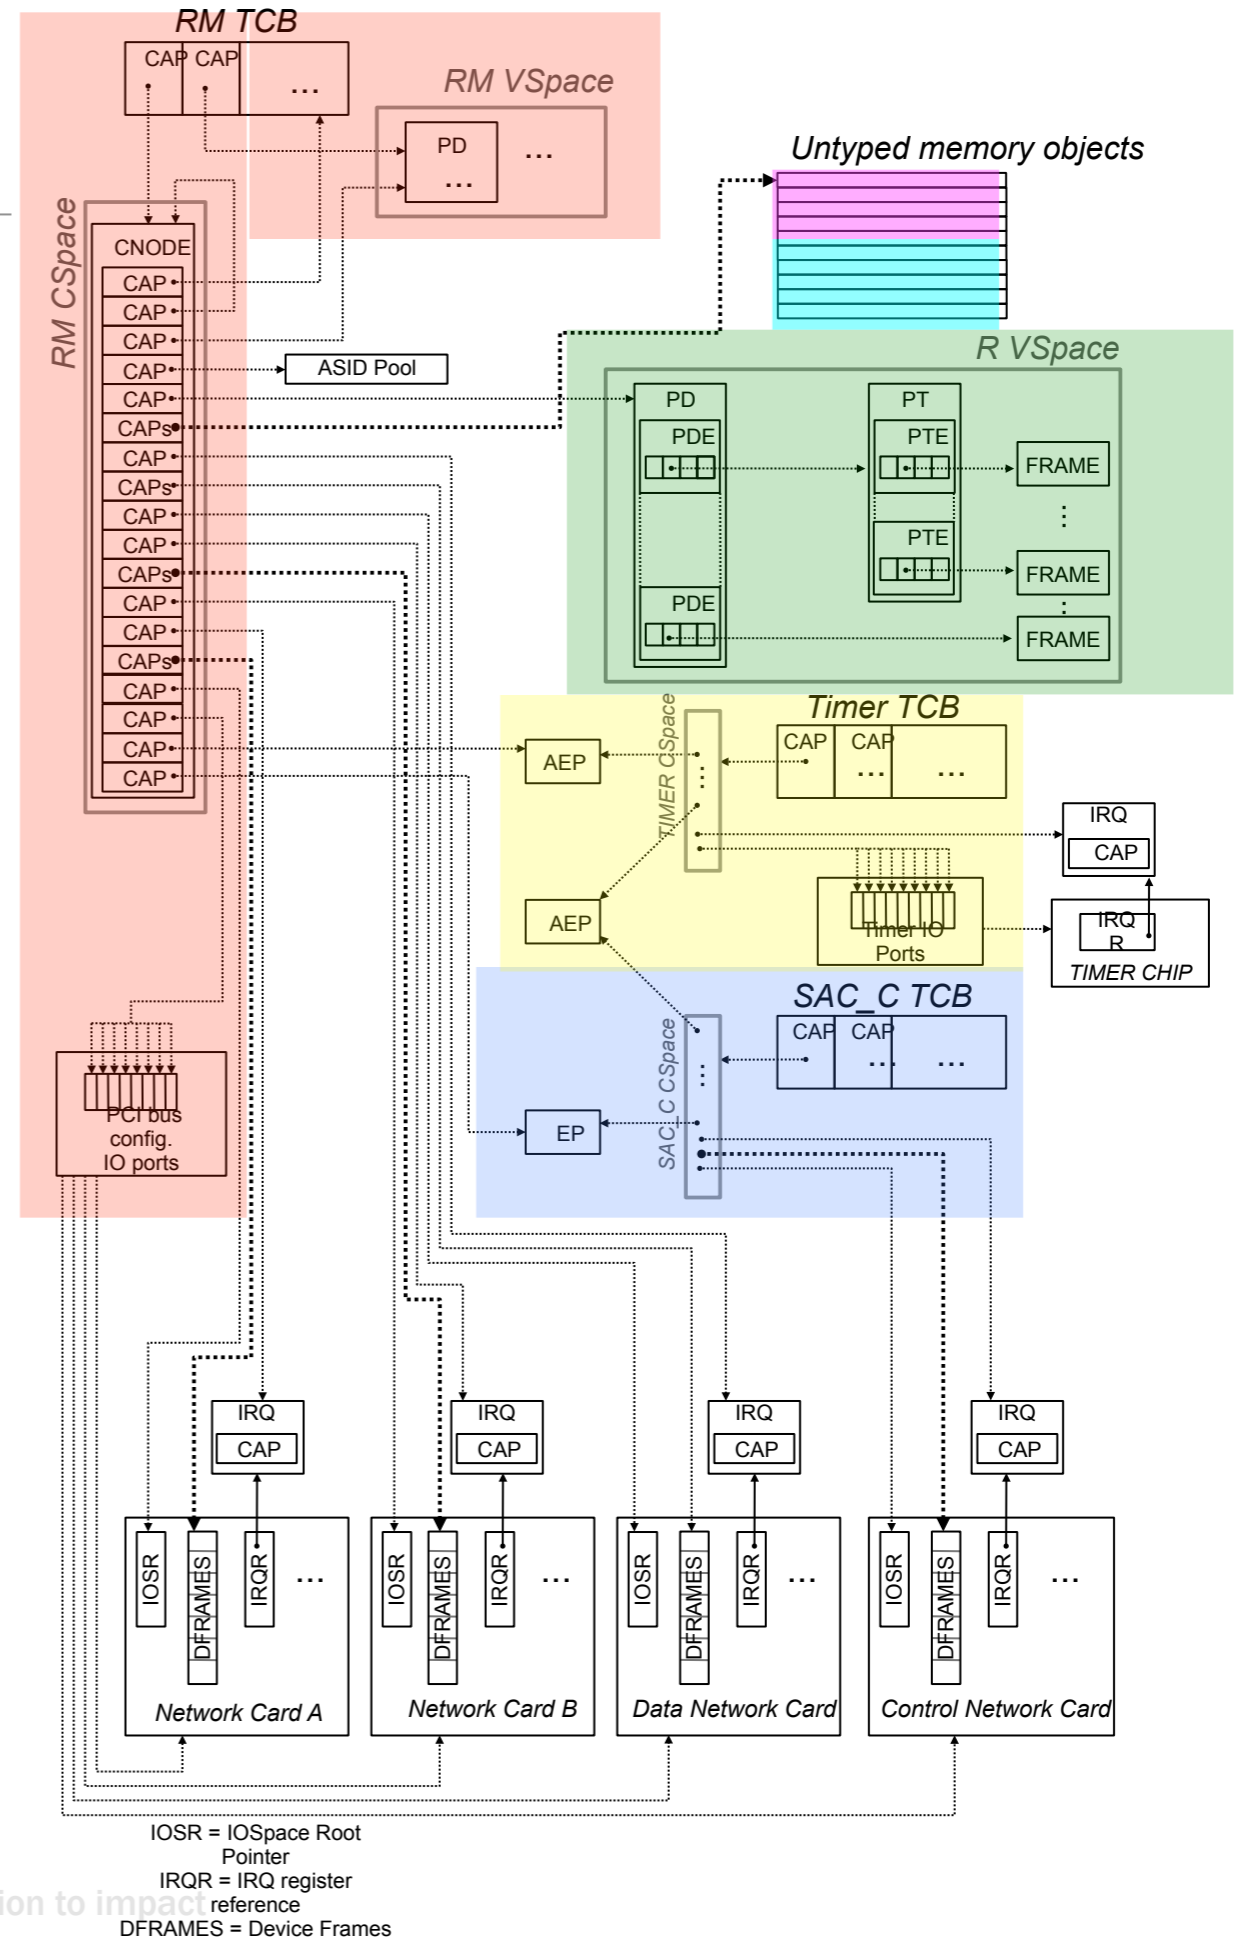
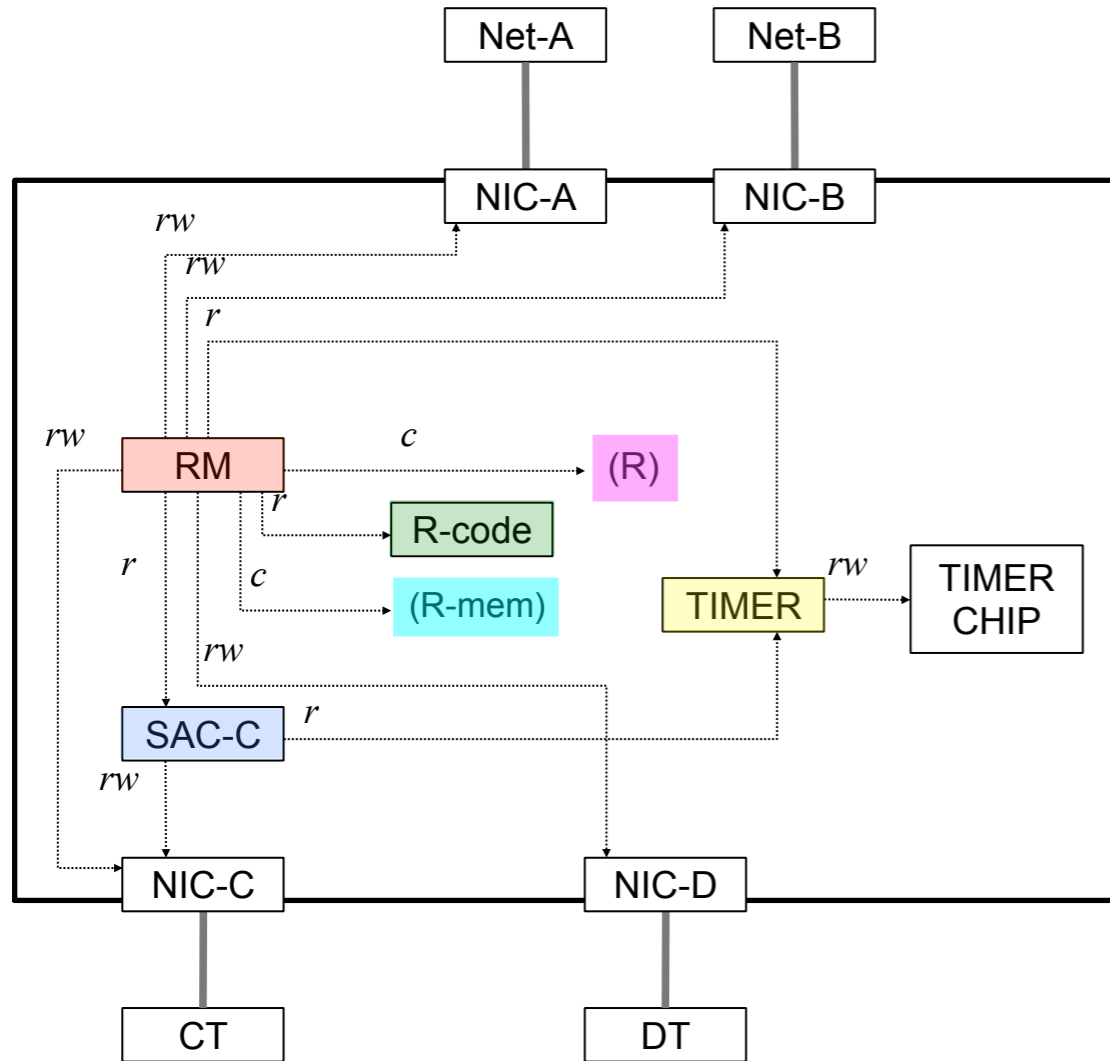


# Low-Level Design

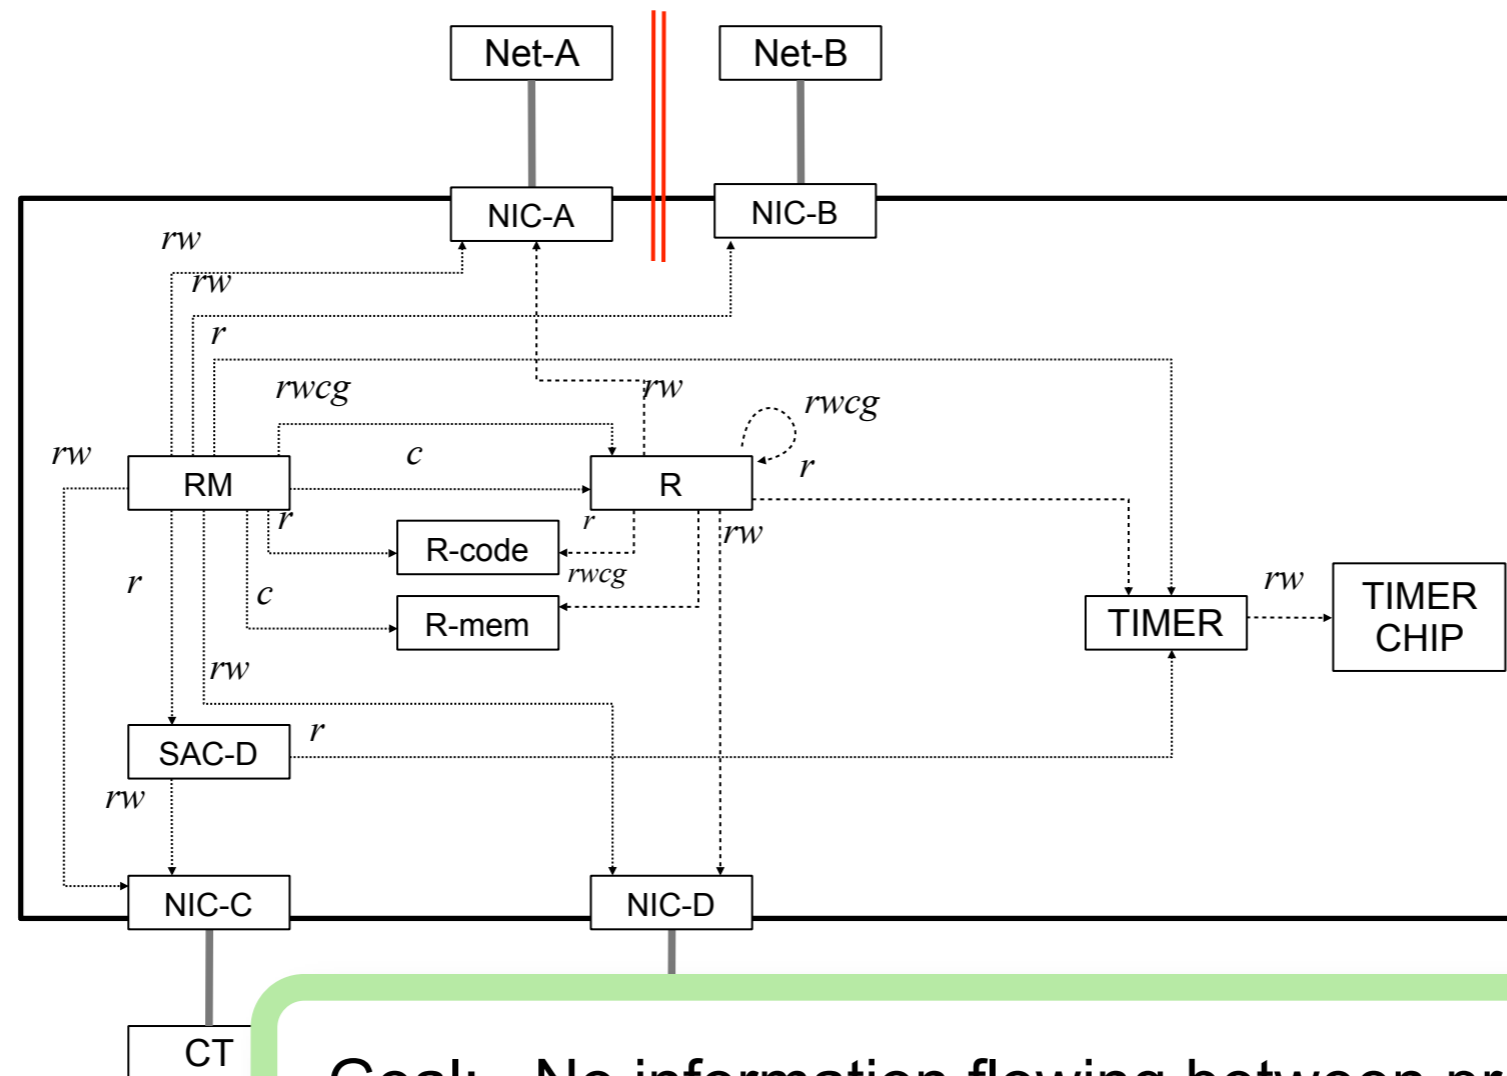


IOSR = IOSpace Root Pointer  
 IRQR = IRQ register reference  
 DFRAMES = Device Frames

# Abstraction

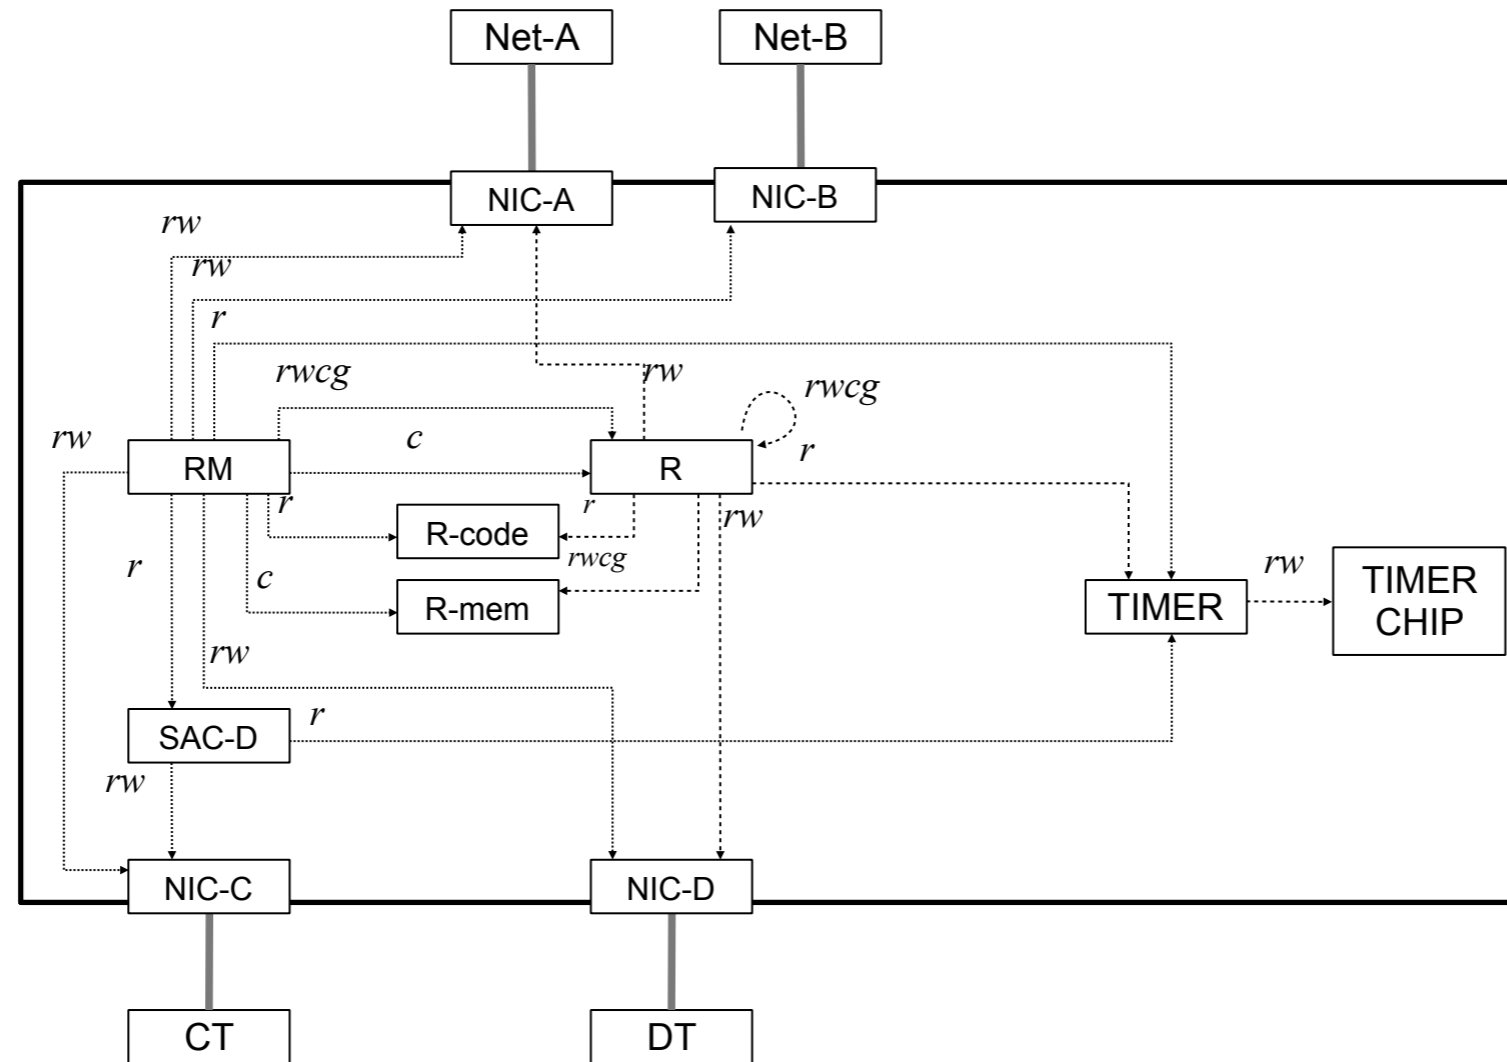


# Security Goal

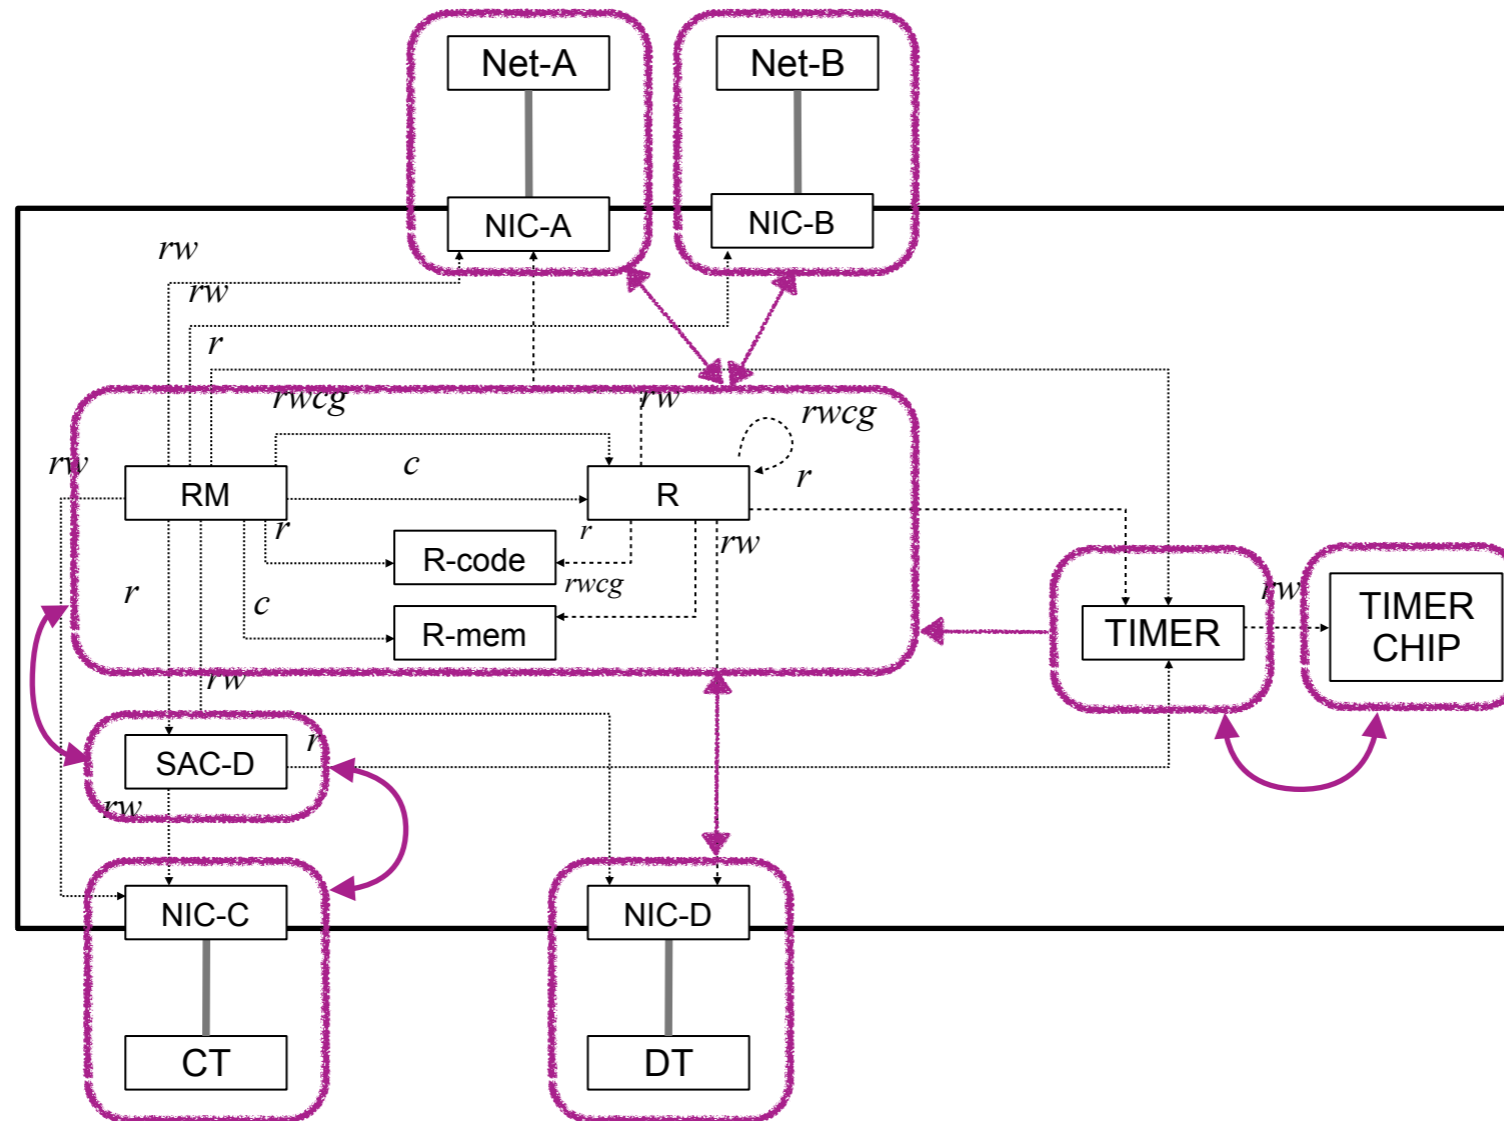


Goal: No information flowing between providers A and B  
Assumption: Info flow through front-end terminal is trusted

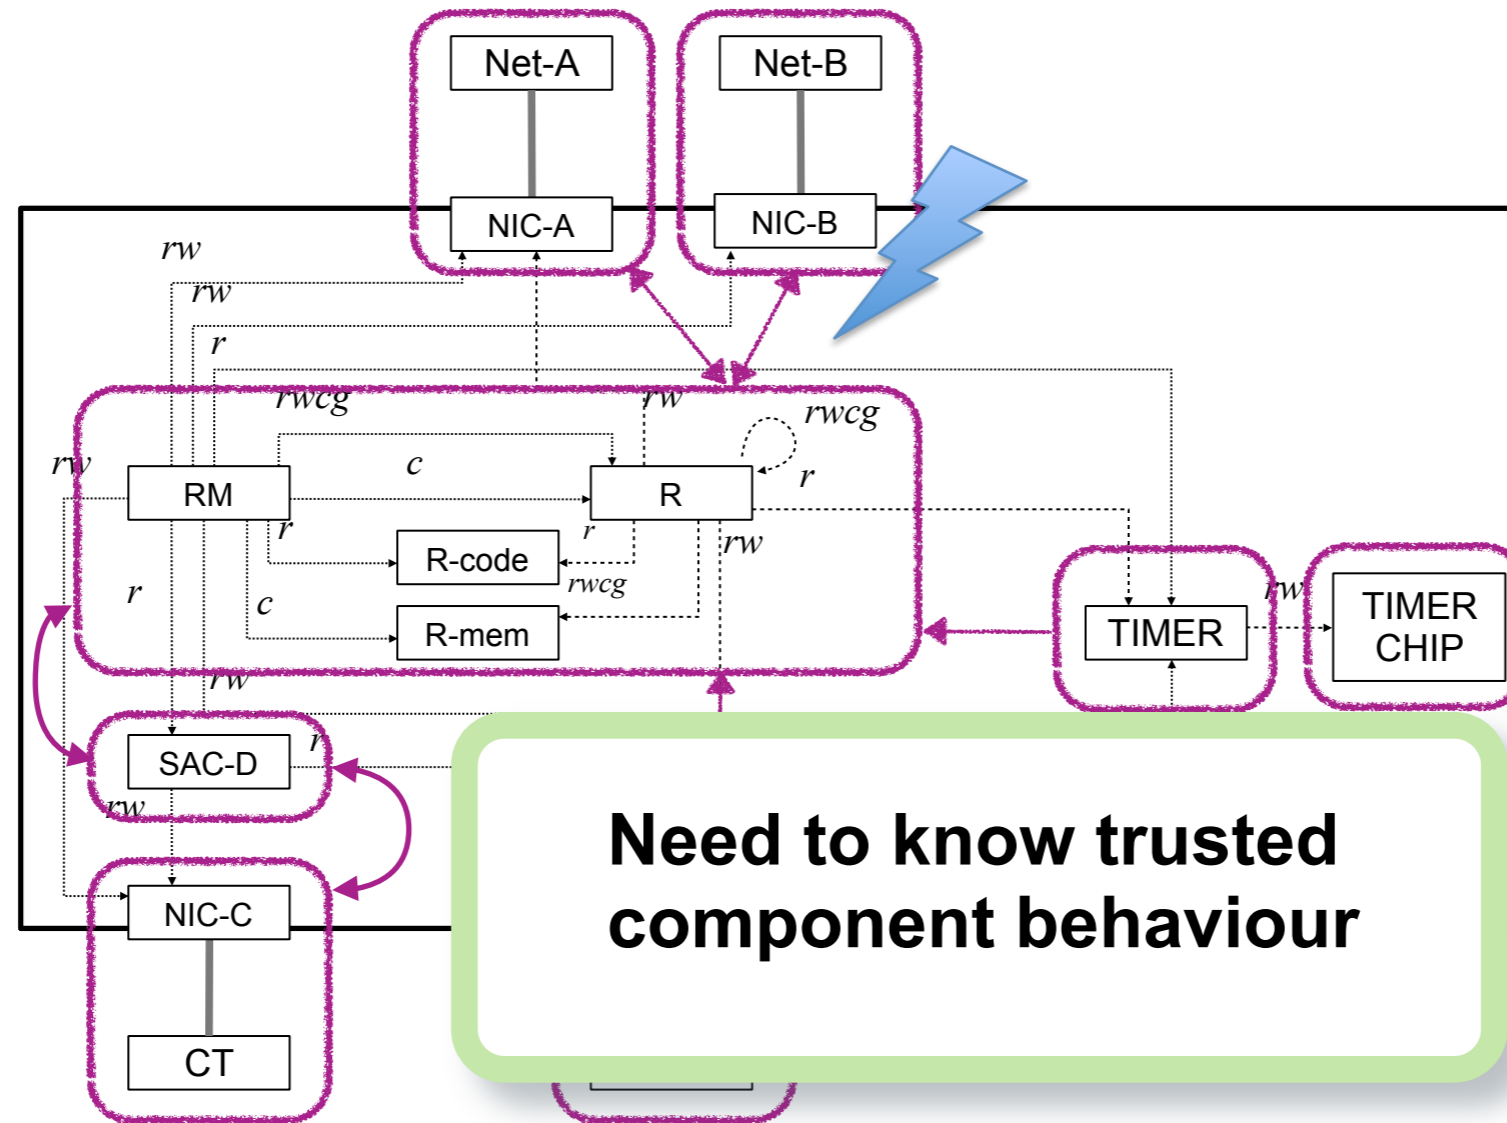
# Plain Take-Grant Analysis



# Plain Take-Grant Analysis

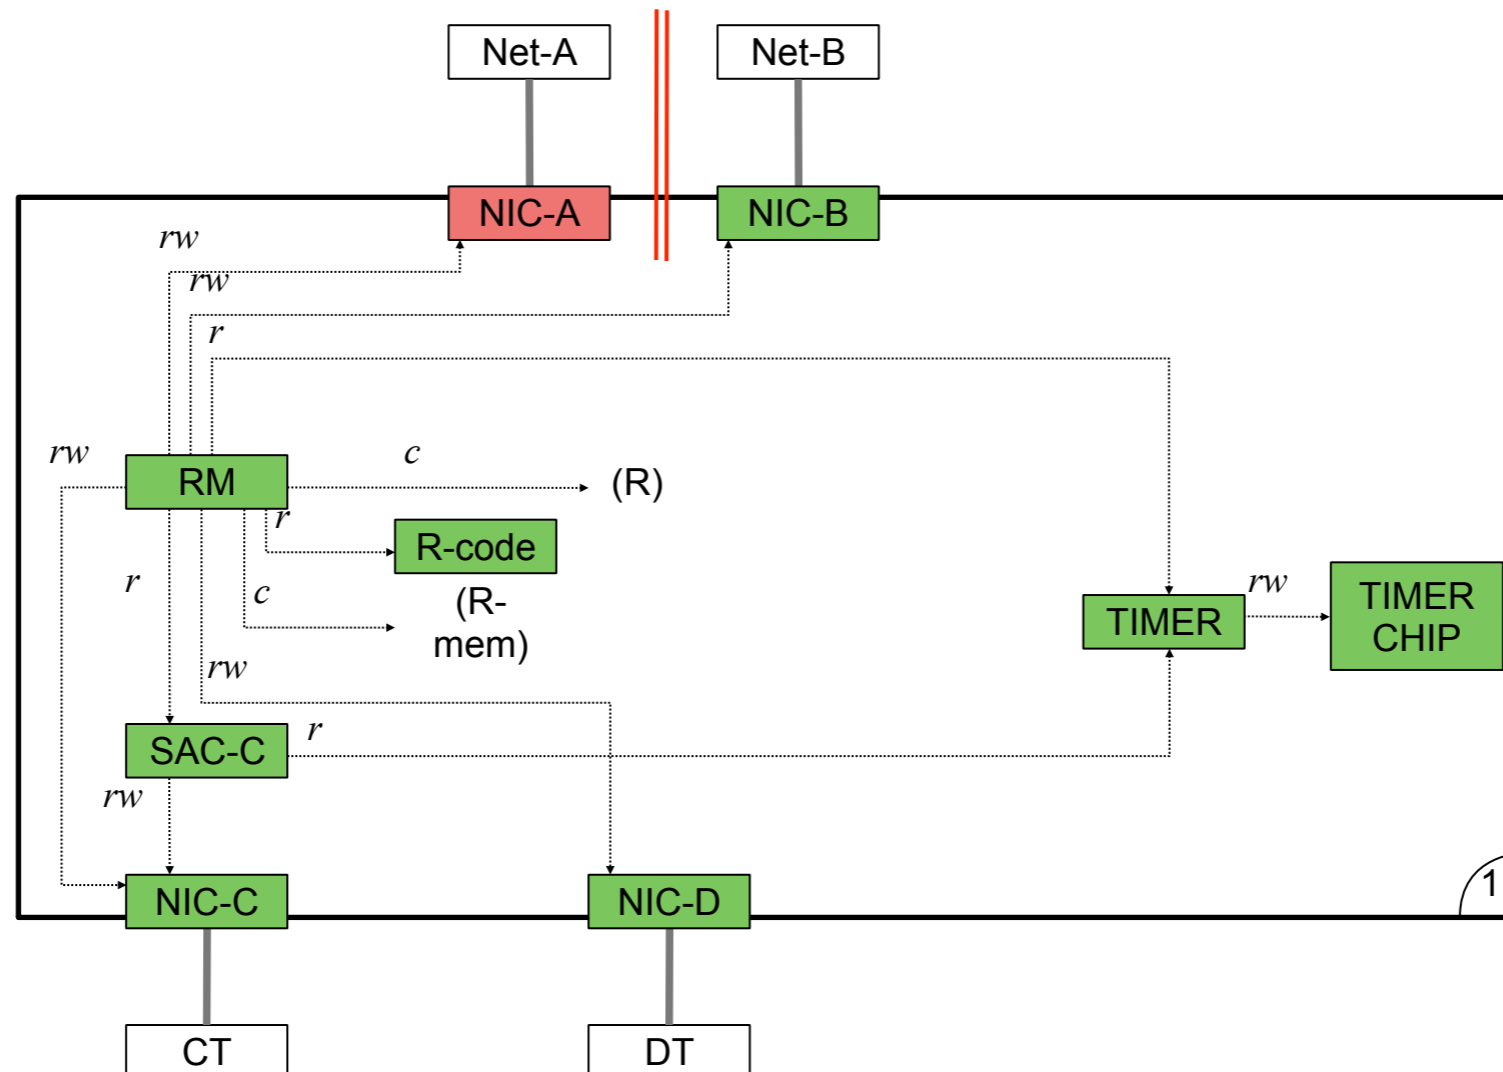


# Plain Take-Grant Analysis





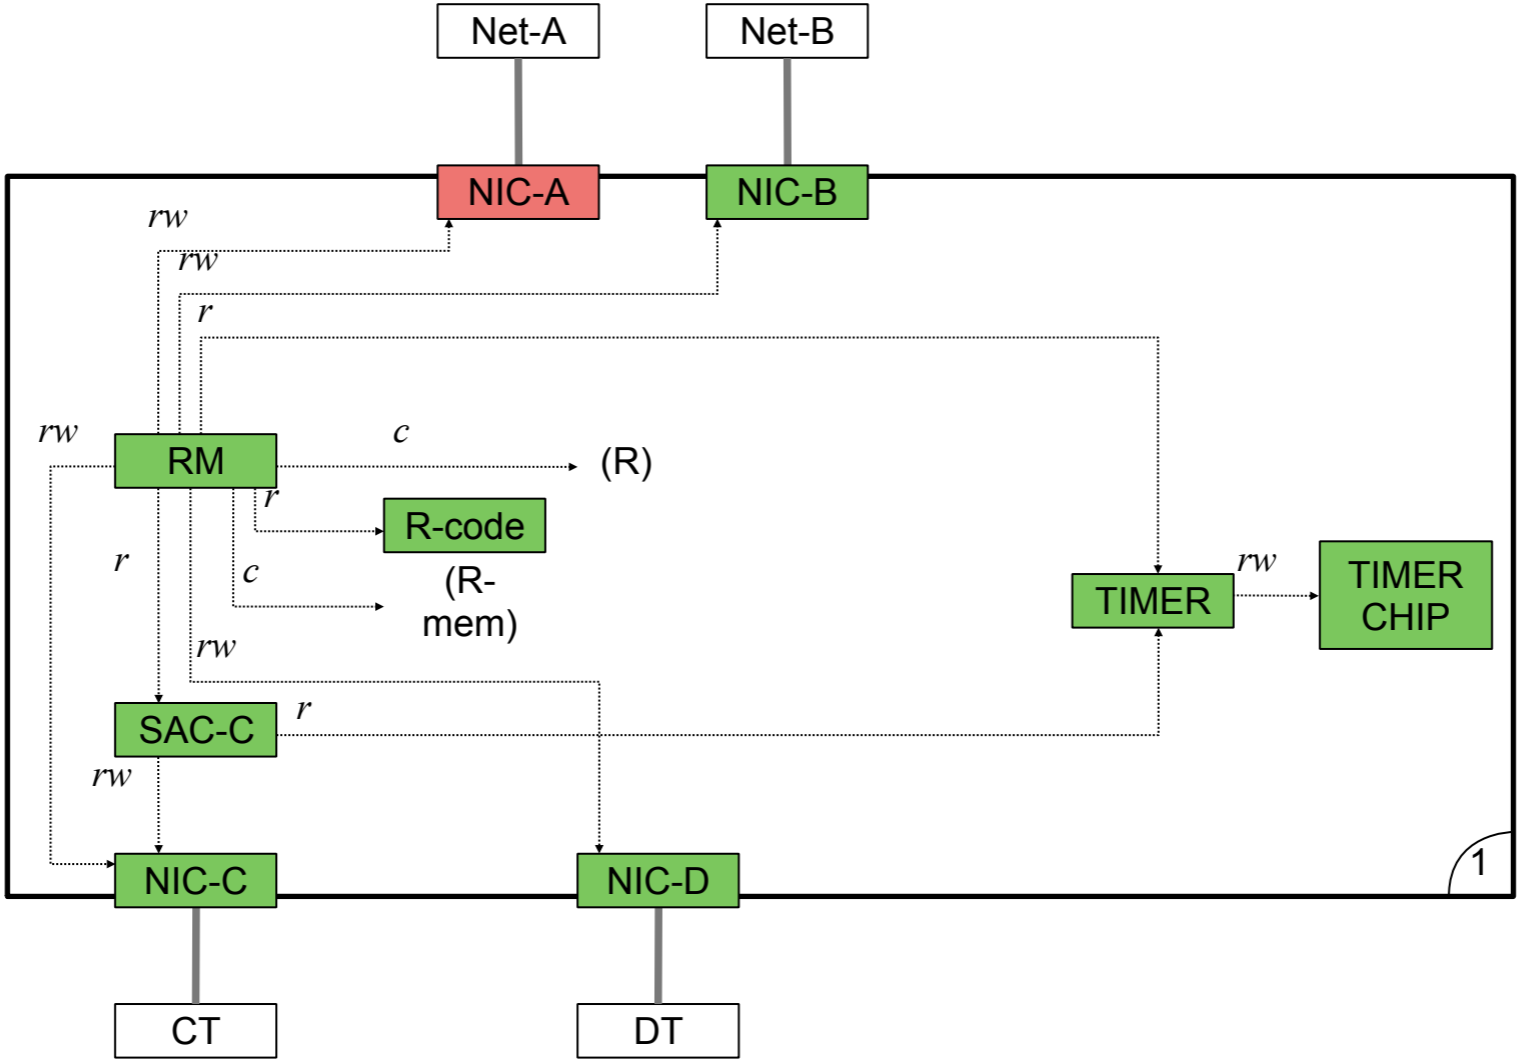
# Security Goal



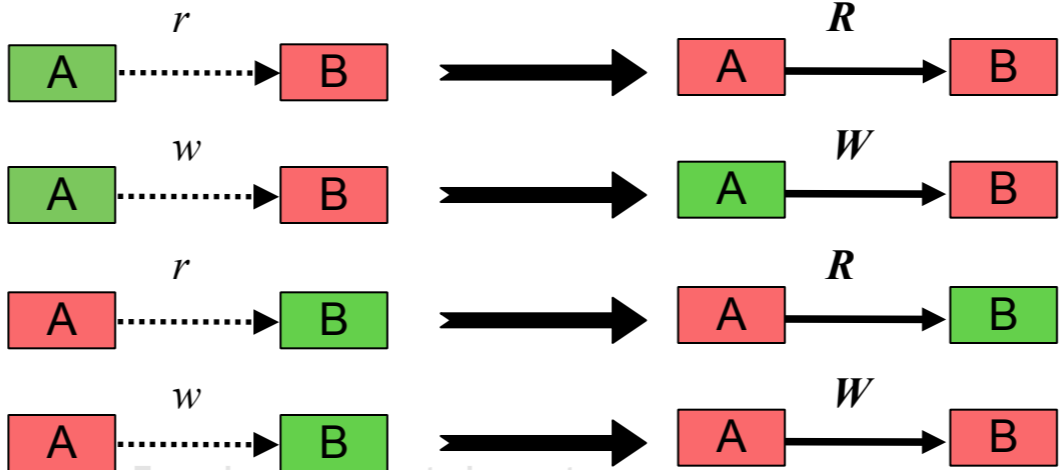
## Approach:

- label-based security
  - tag as 'contaminated' if may contain data from Net-A
  - NIC-A always contaminated
- Goal: prove NIC-B always 'not contaminated'


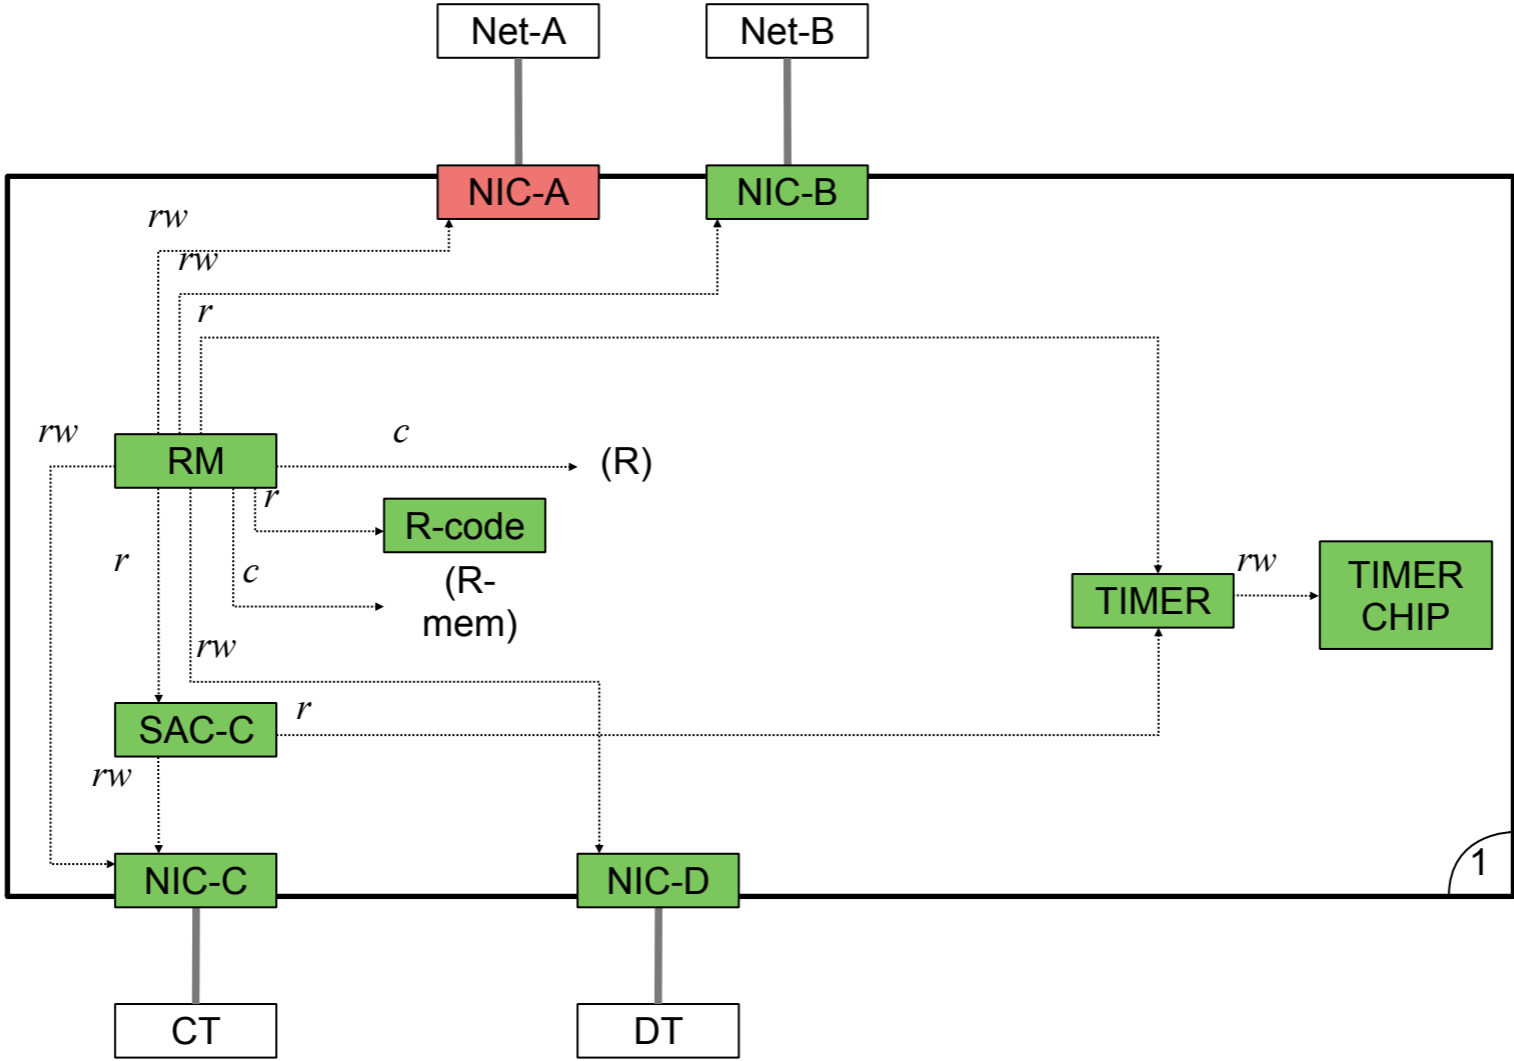
# Security Analysis



Rules:




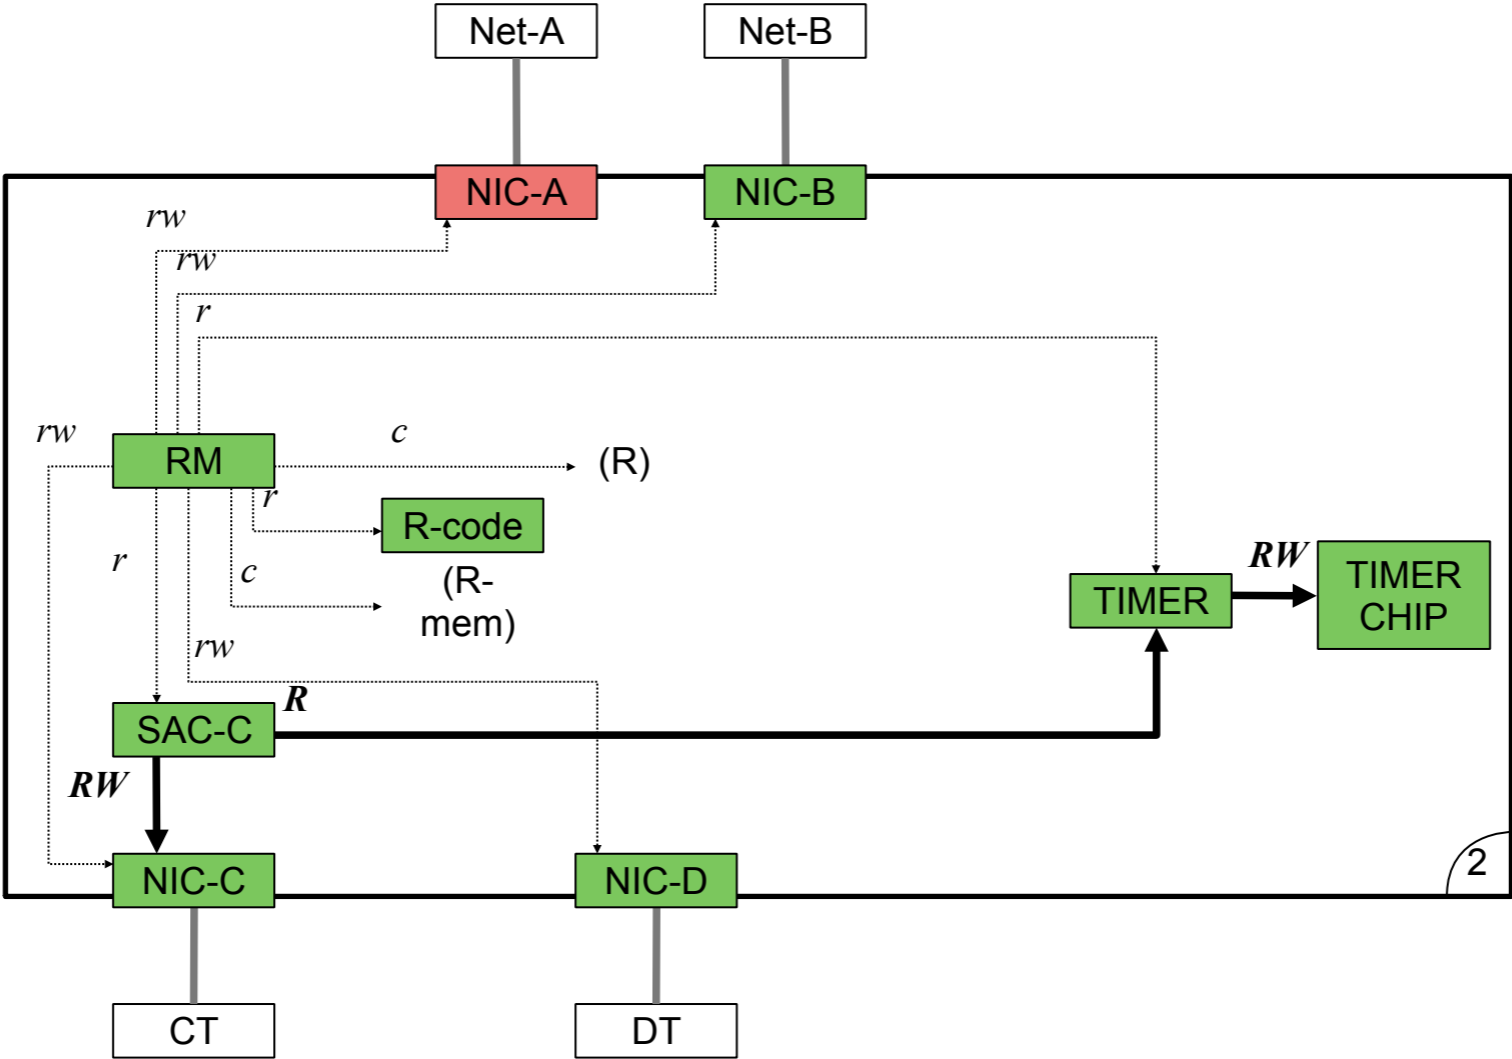
# Life Cycle



# Life Cycle




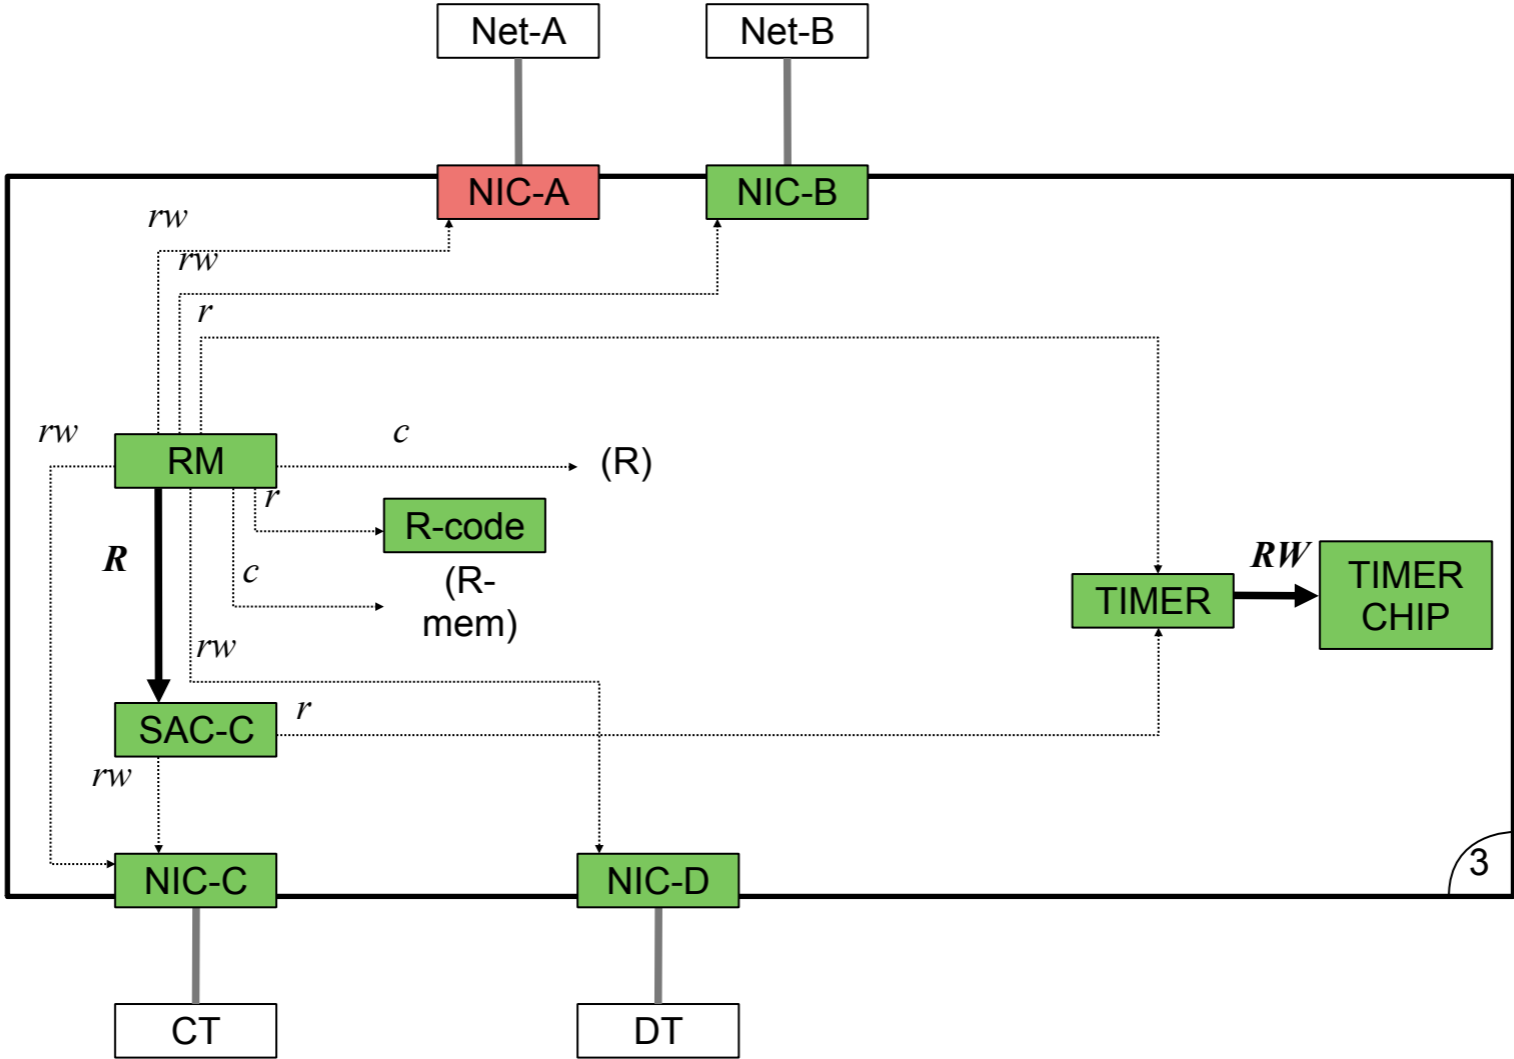
CT authenticates with SAC-C  
**→**  
 CT sends a request to switch to Net-A



# Life Cycle



RM receives request to switch to Net-A


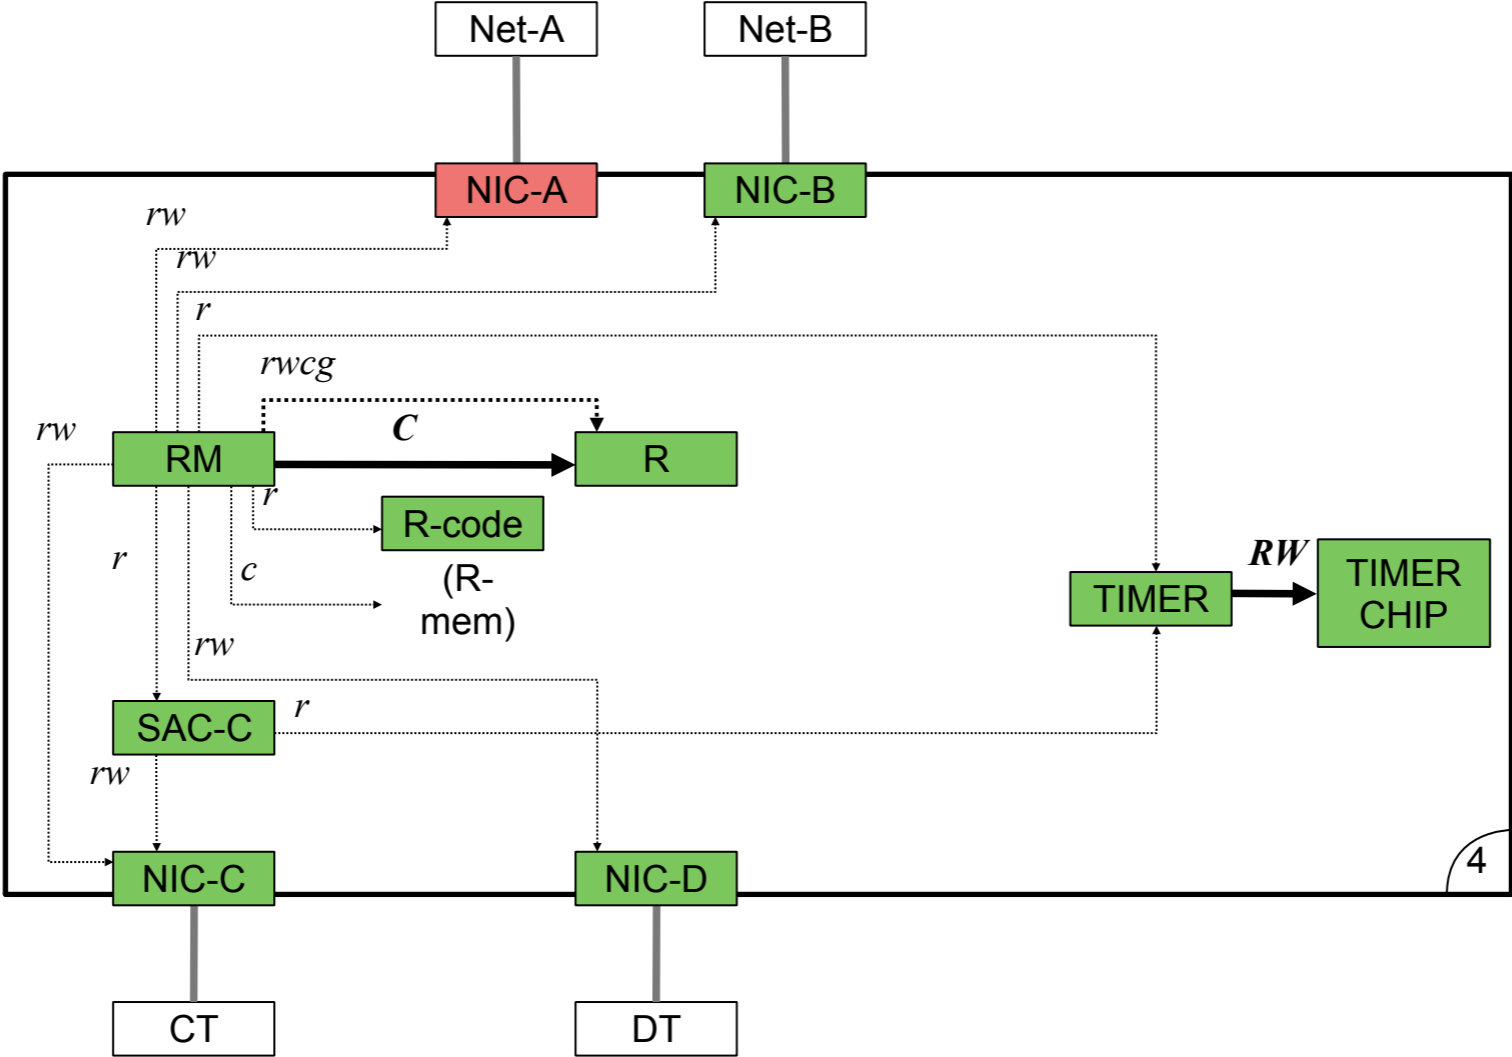


3

# Life Cycle

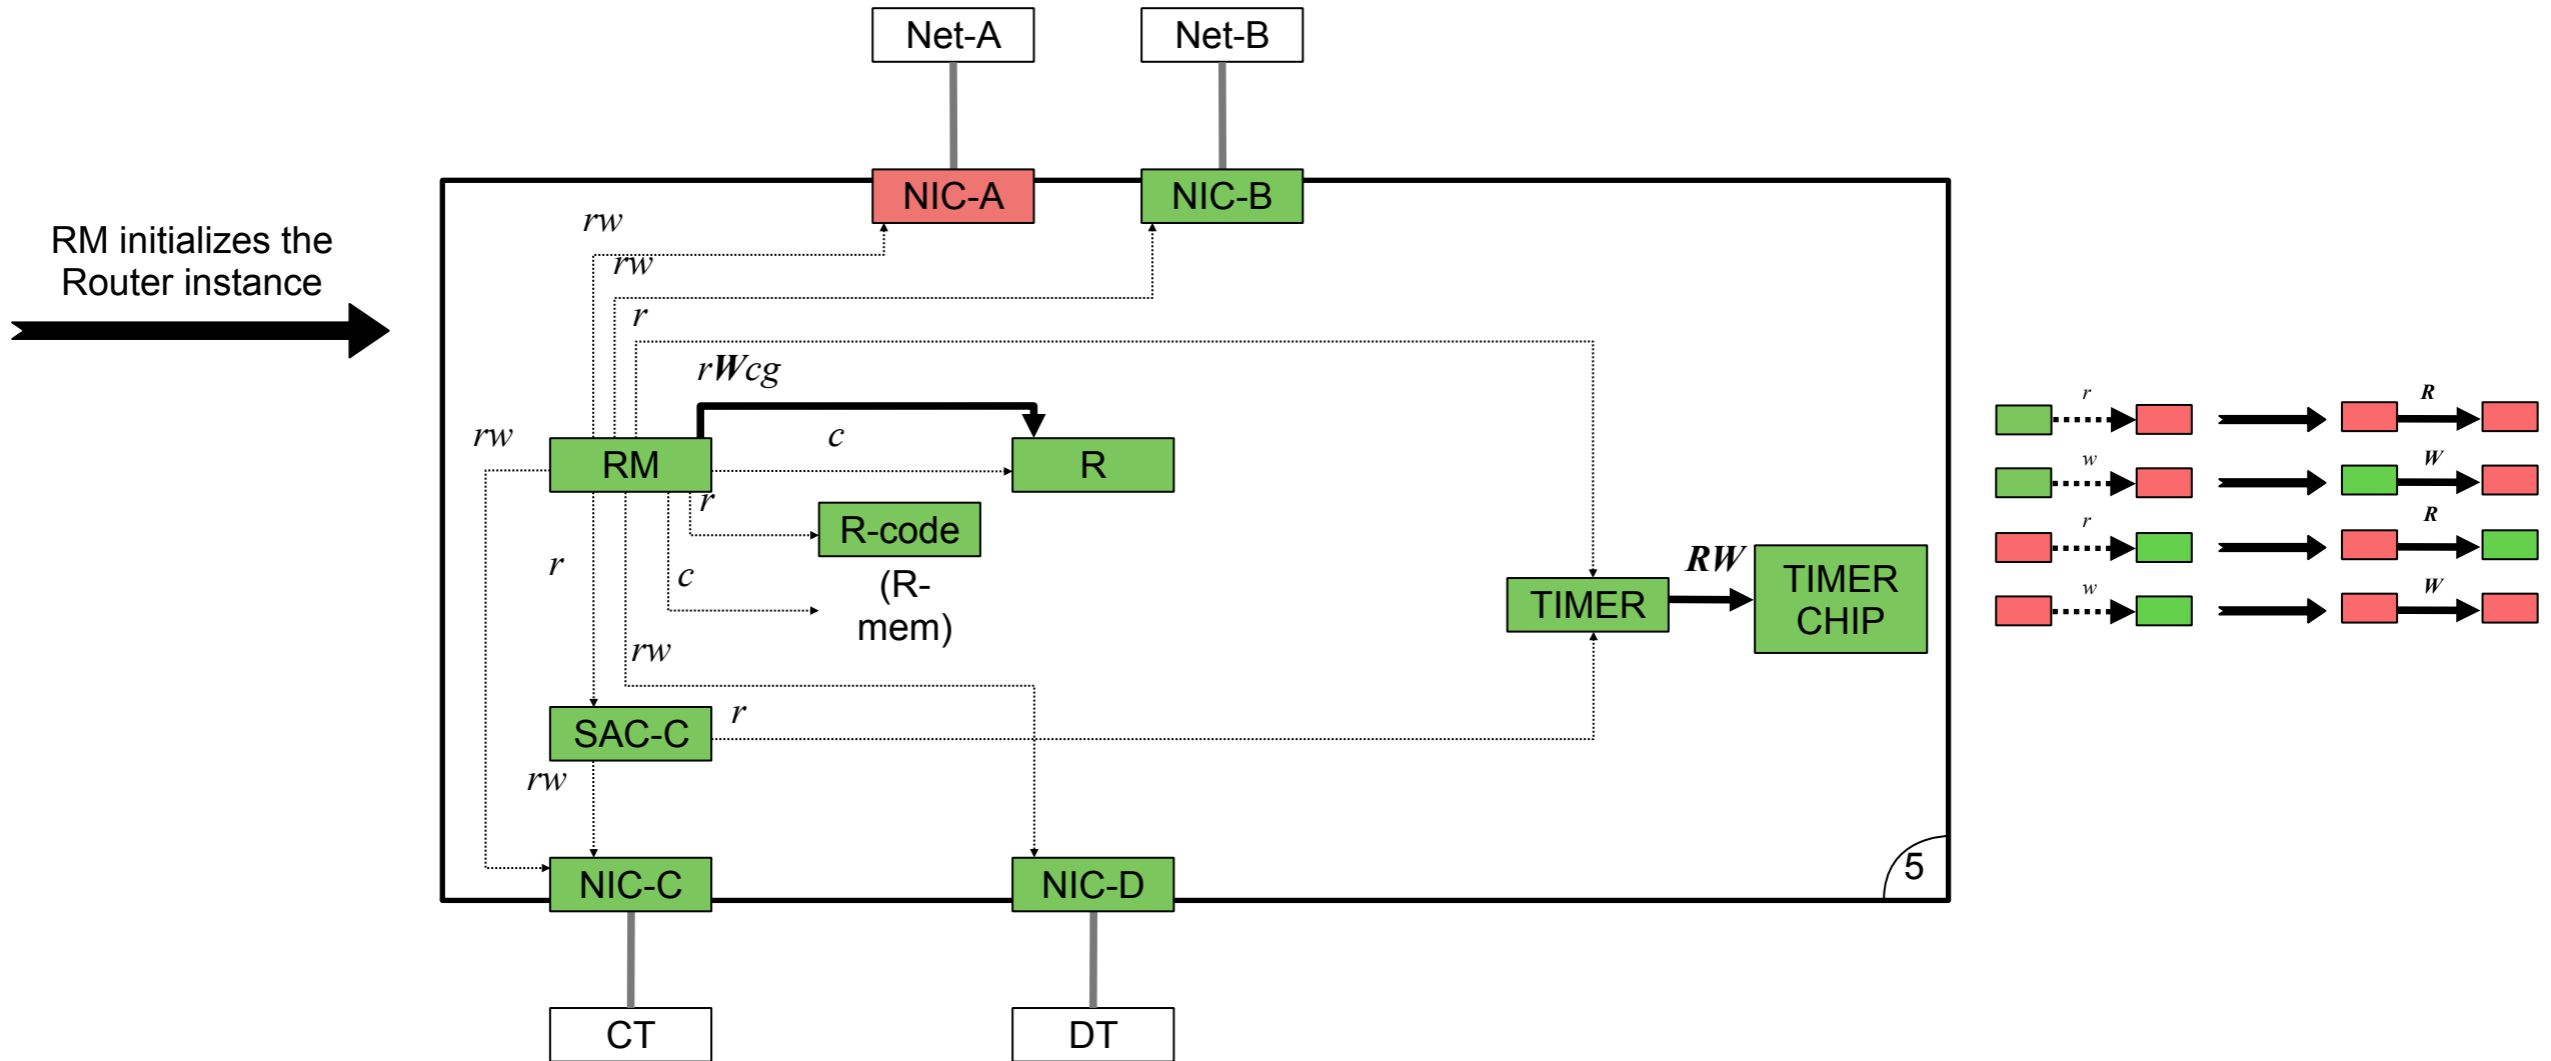


RM creates a new Router instance R  
 (and gets full rights to the newly created object)






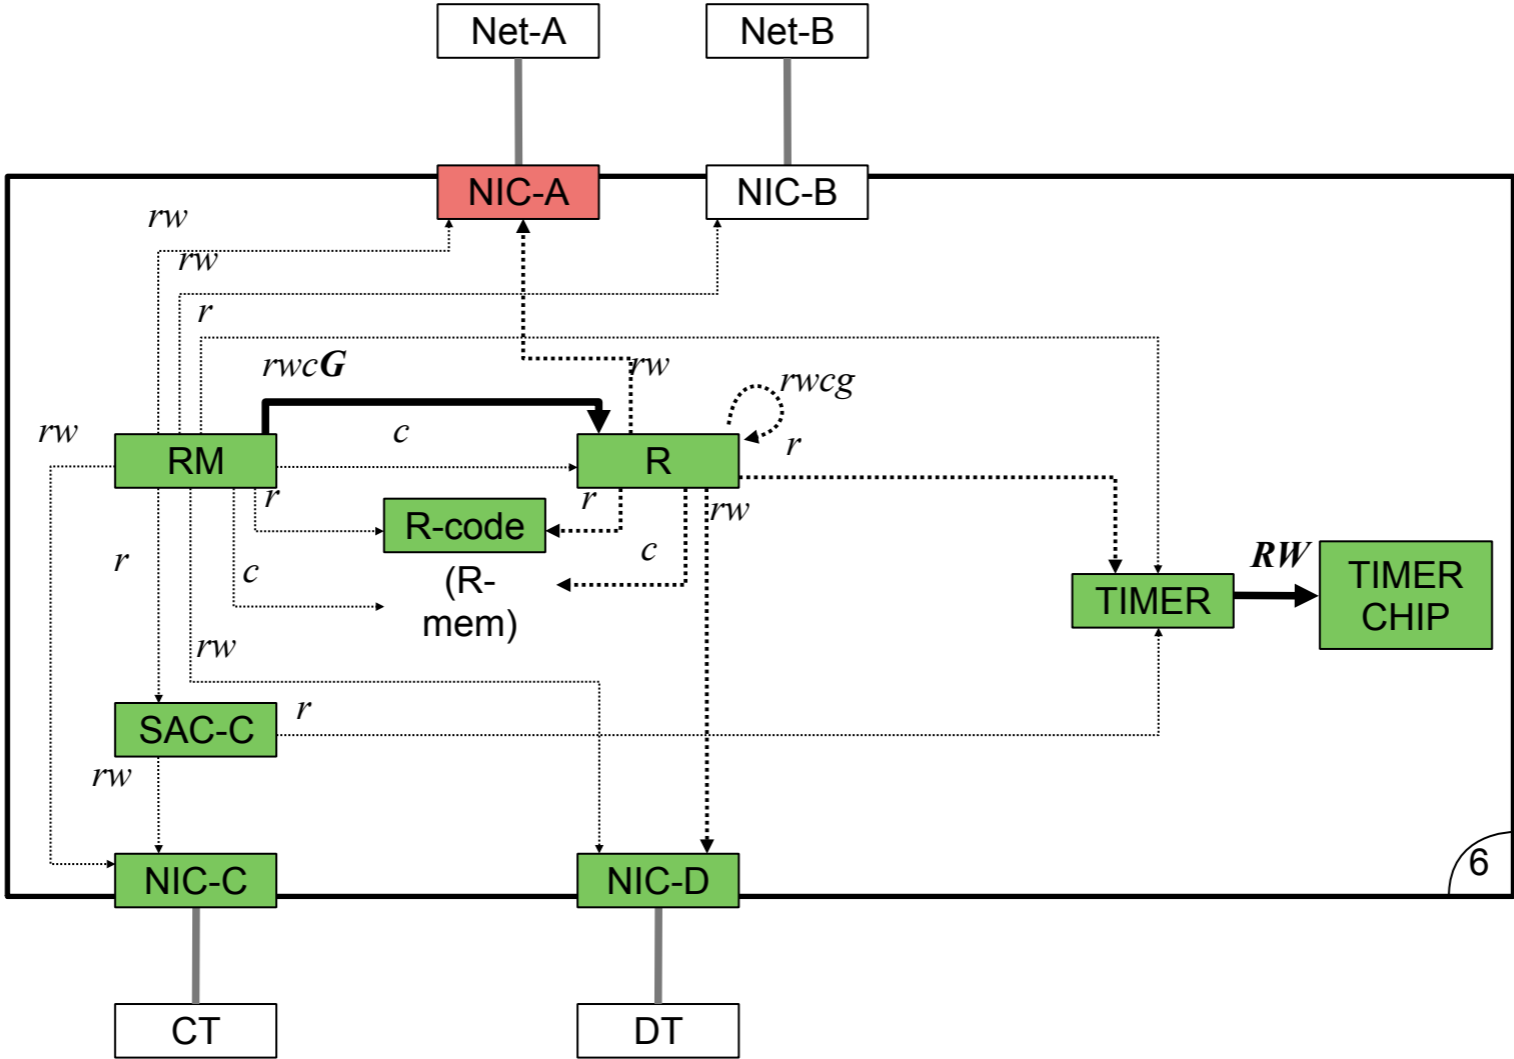
# Life Cycle



# Life Cycle



RM grants to R its rights to NIC-A, NIC-D, R-mem, R-code, TIMER, and itself




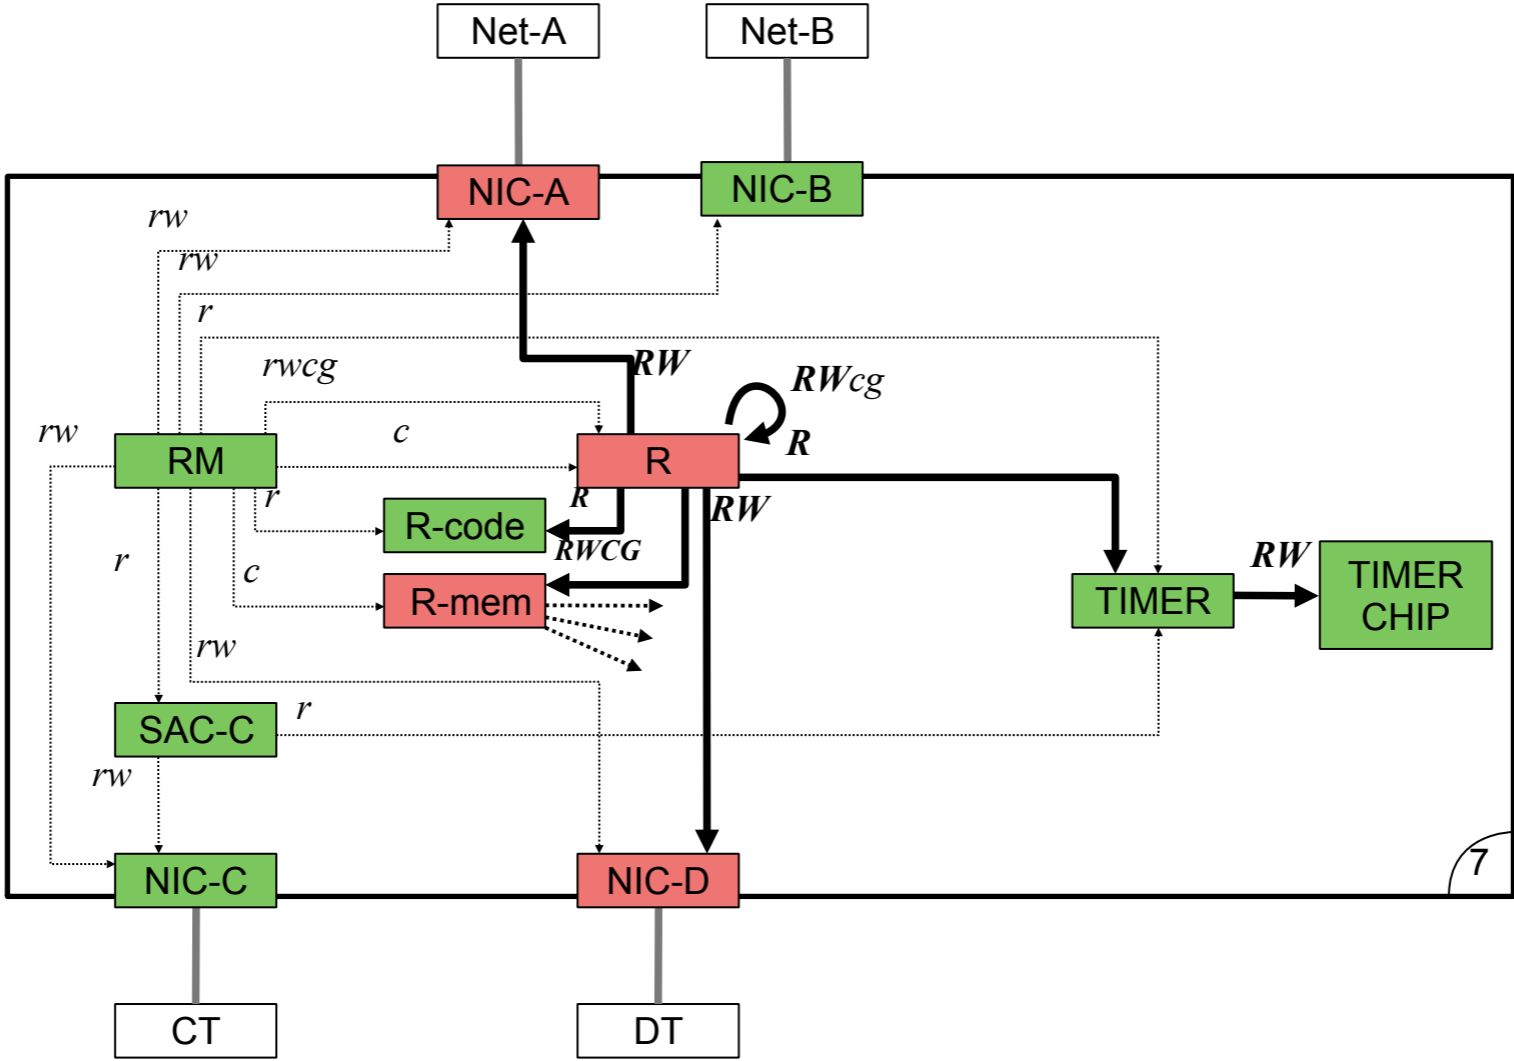
# Life Cycle



DT starts to communicate with Net-A, through R.

**→**


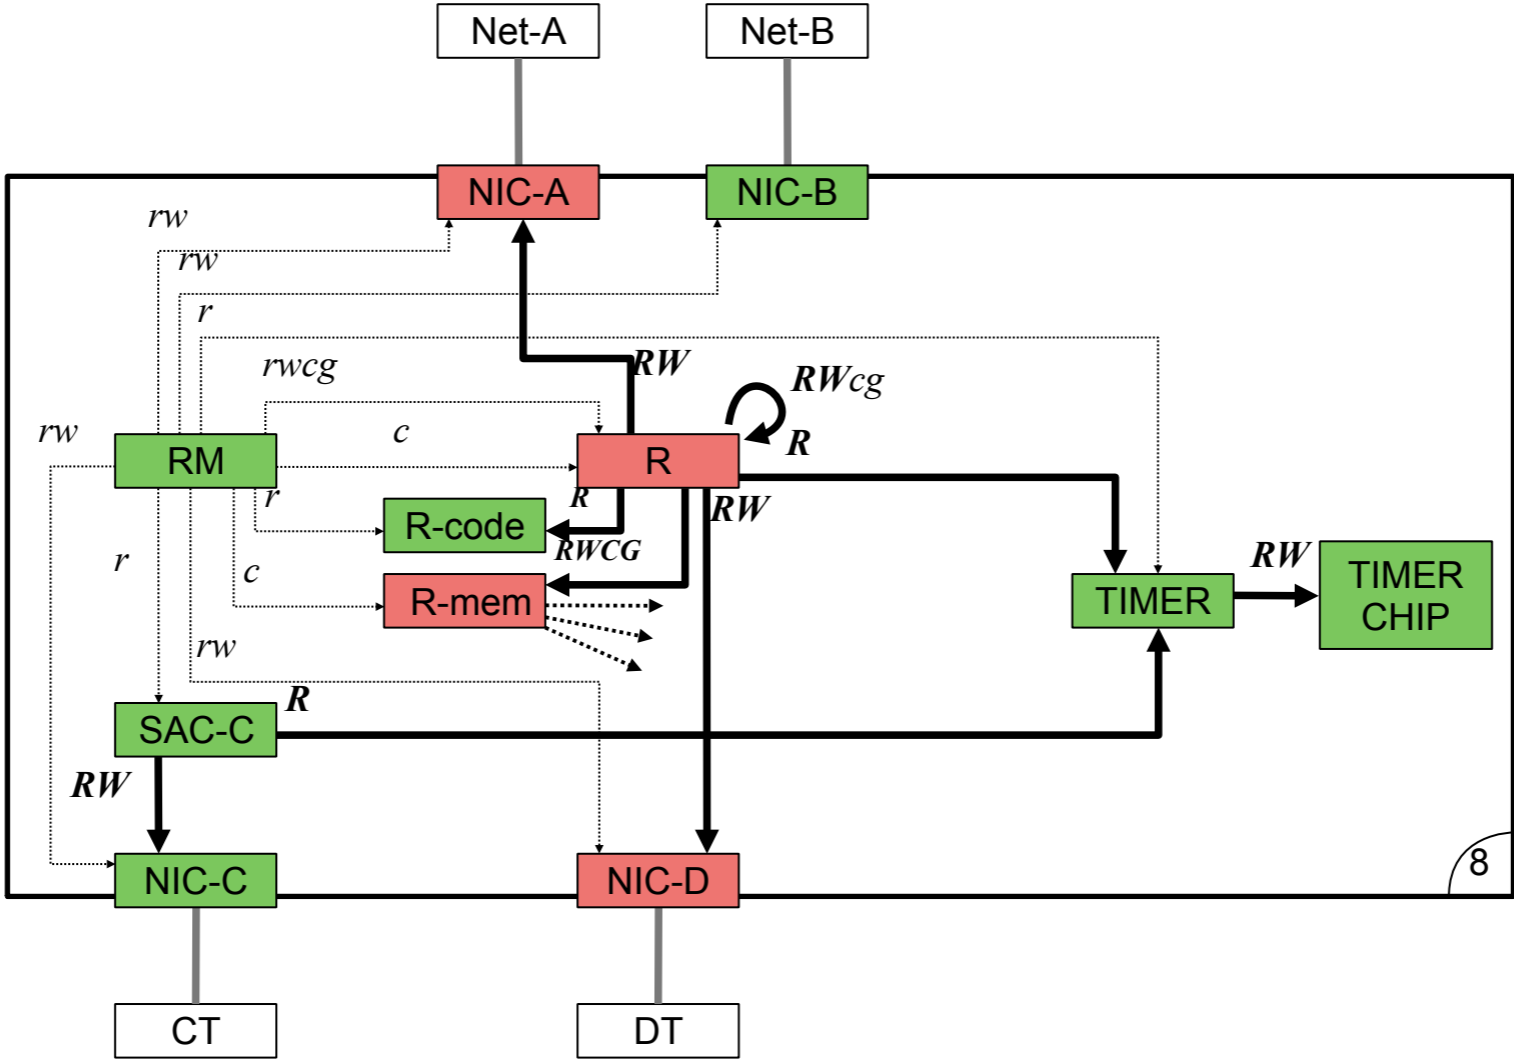
This may imply the creation of new objects using R-mem and granting caps to them.



# Life Cycle




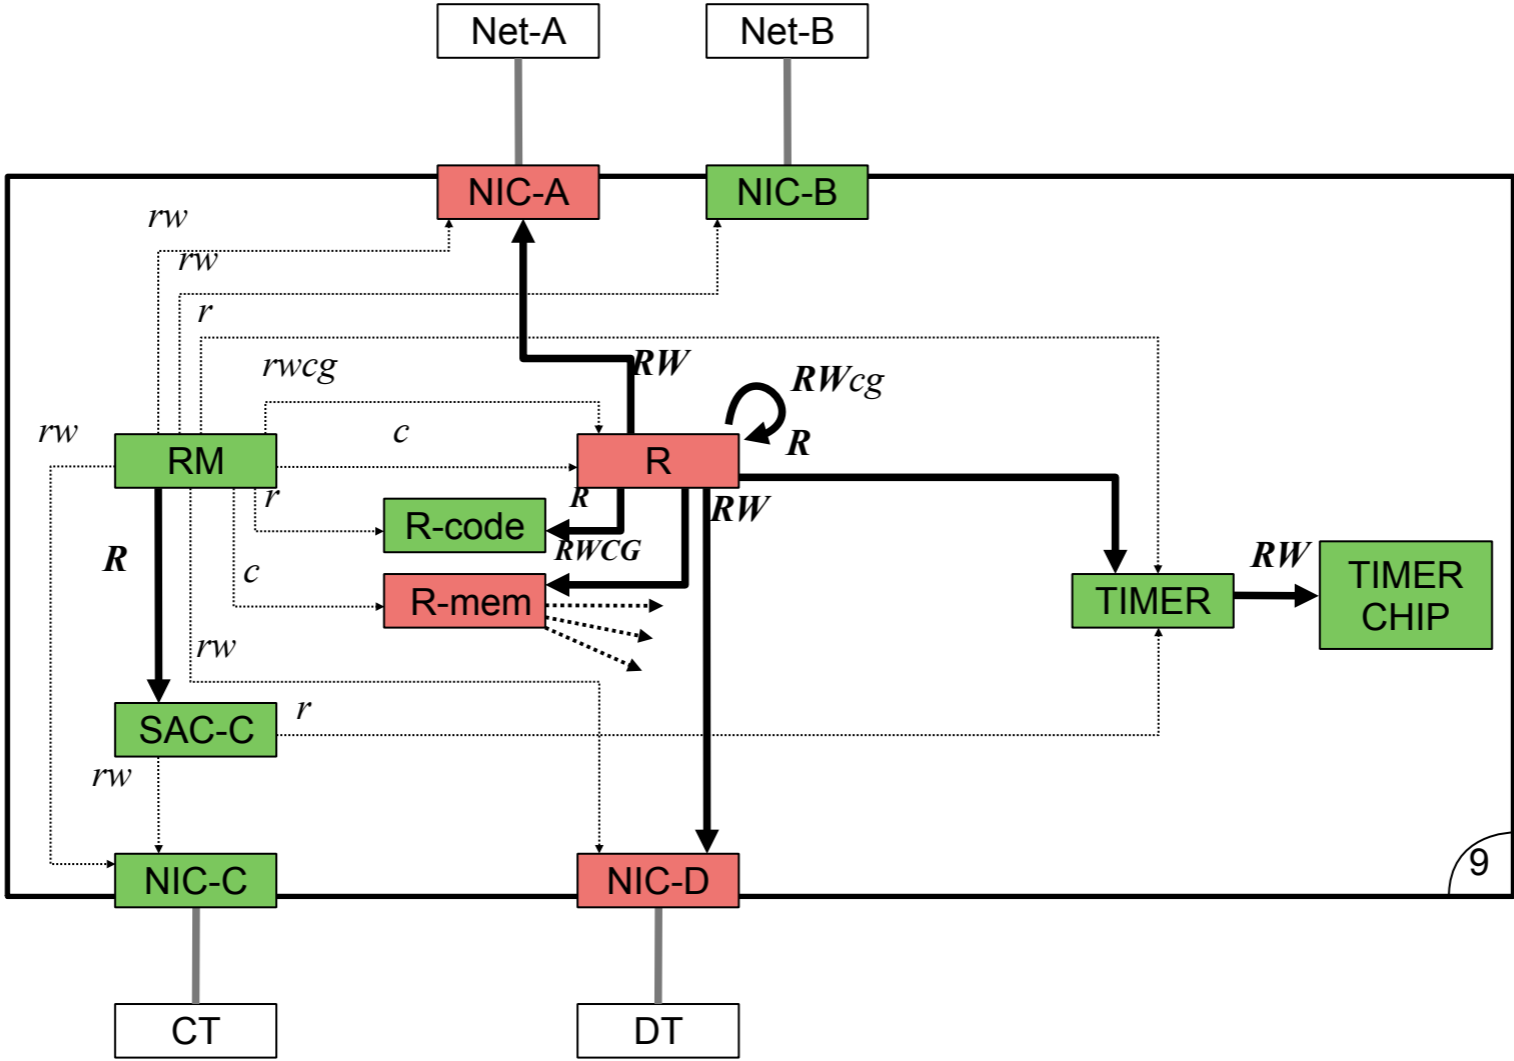
CT sends a request to switch to Net-B  
 (while DT still communicates with Net-A through R)



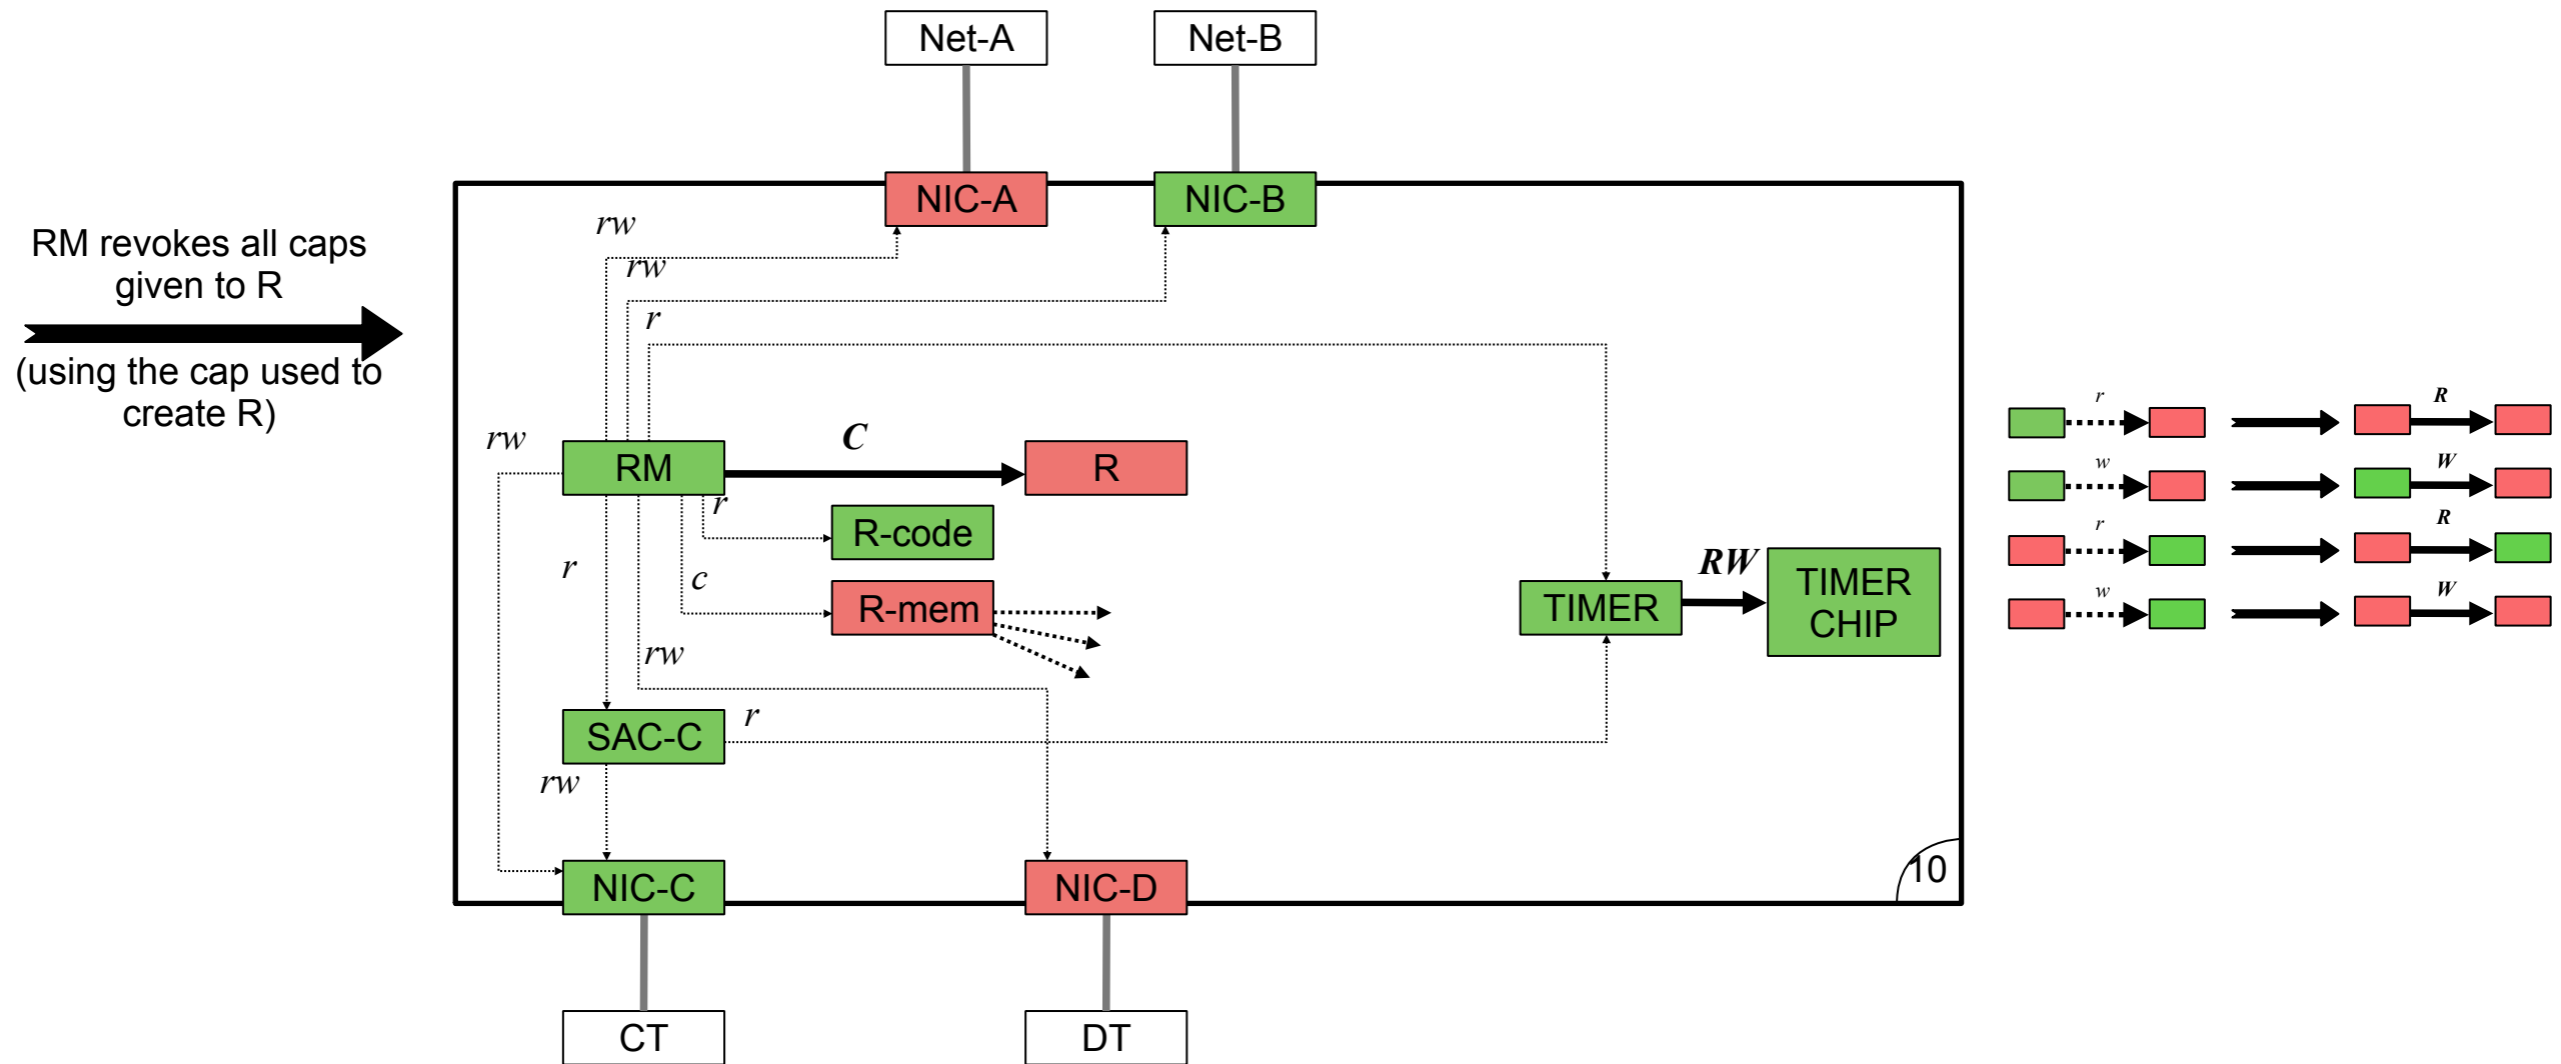
# Life Cycle



RM receives request to switch to Net-B




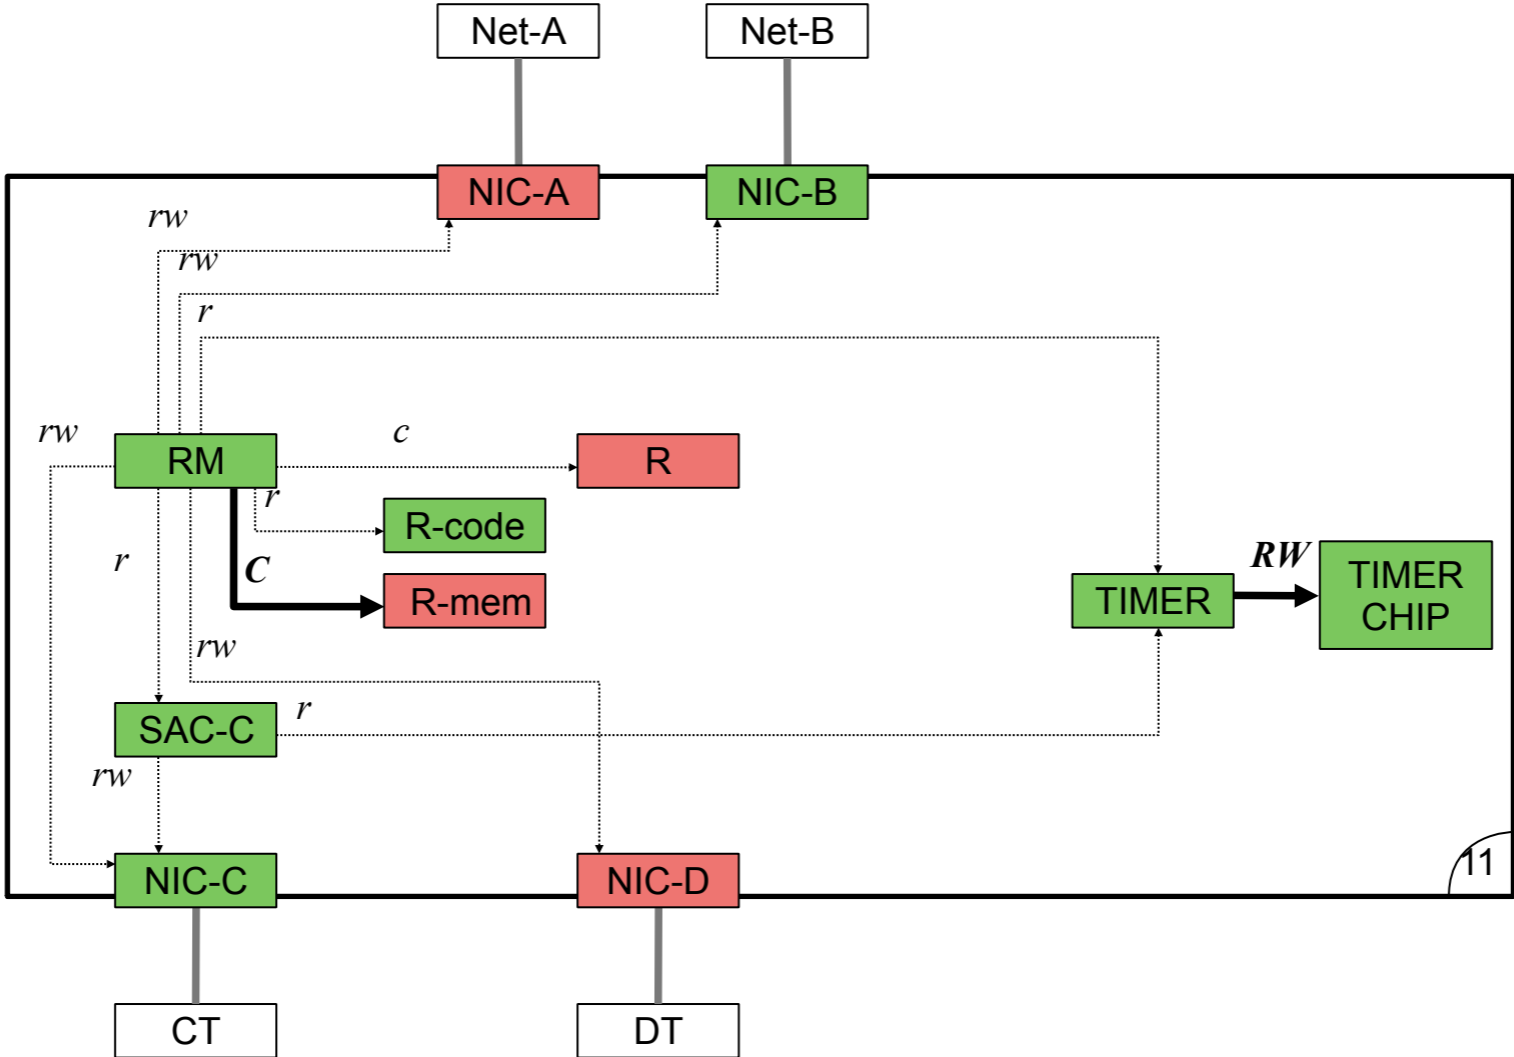
# Life Cycle



# Life Cycle

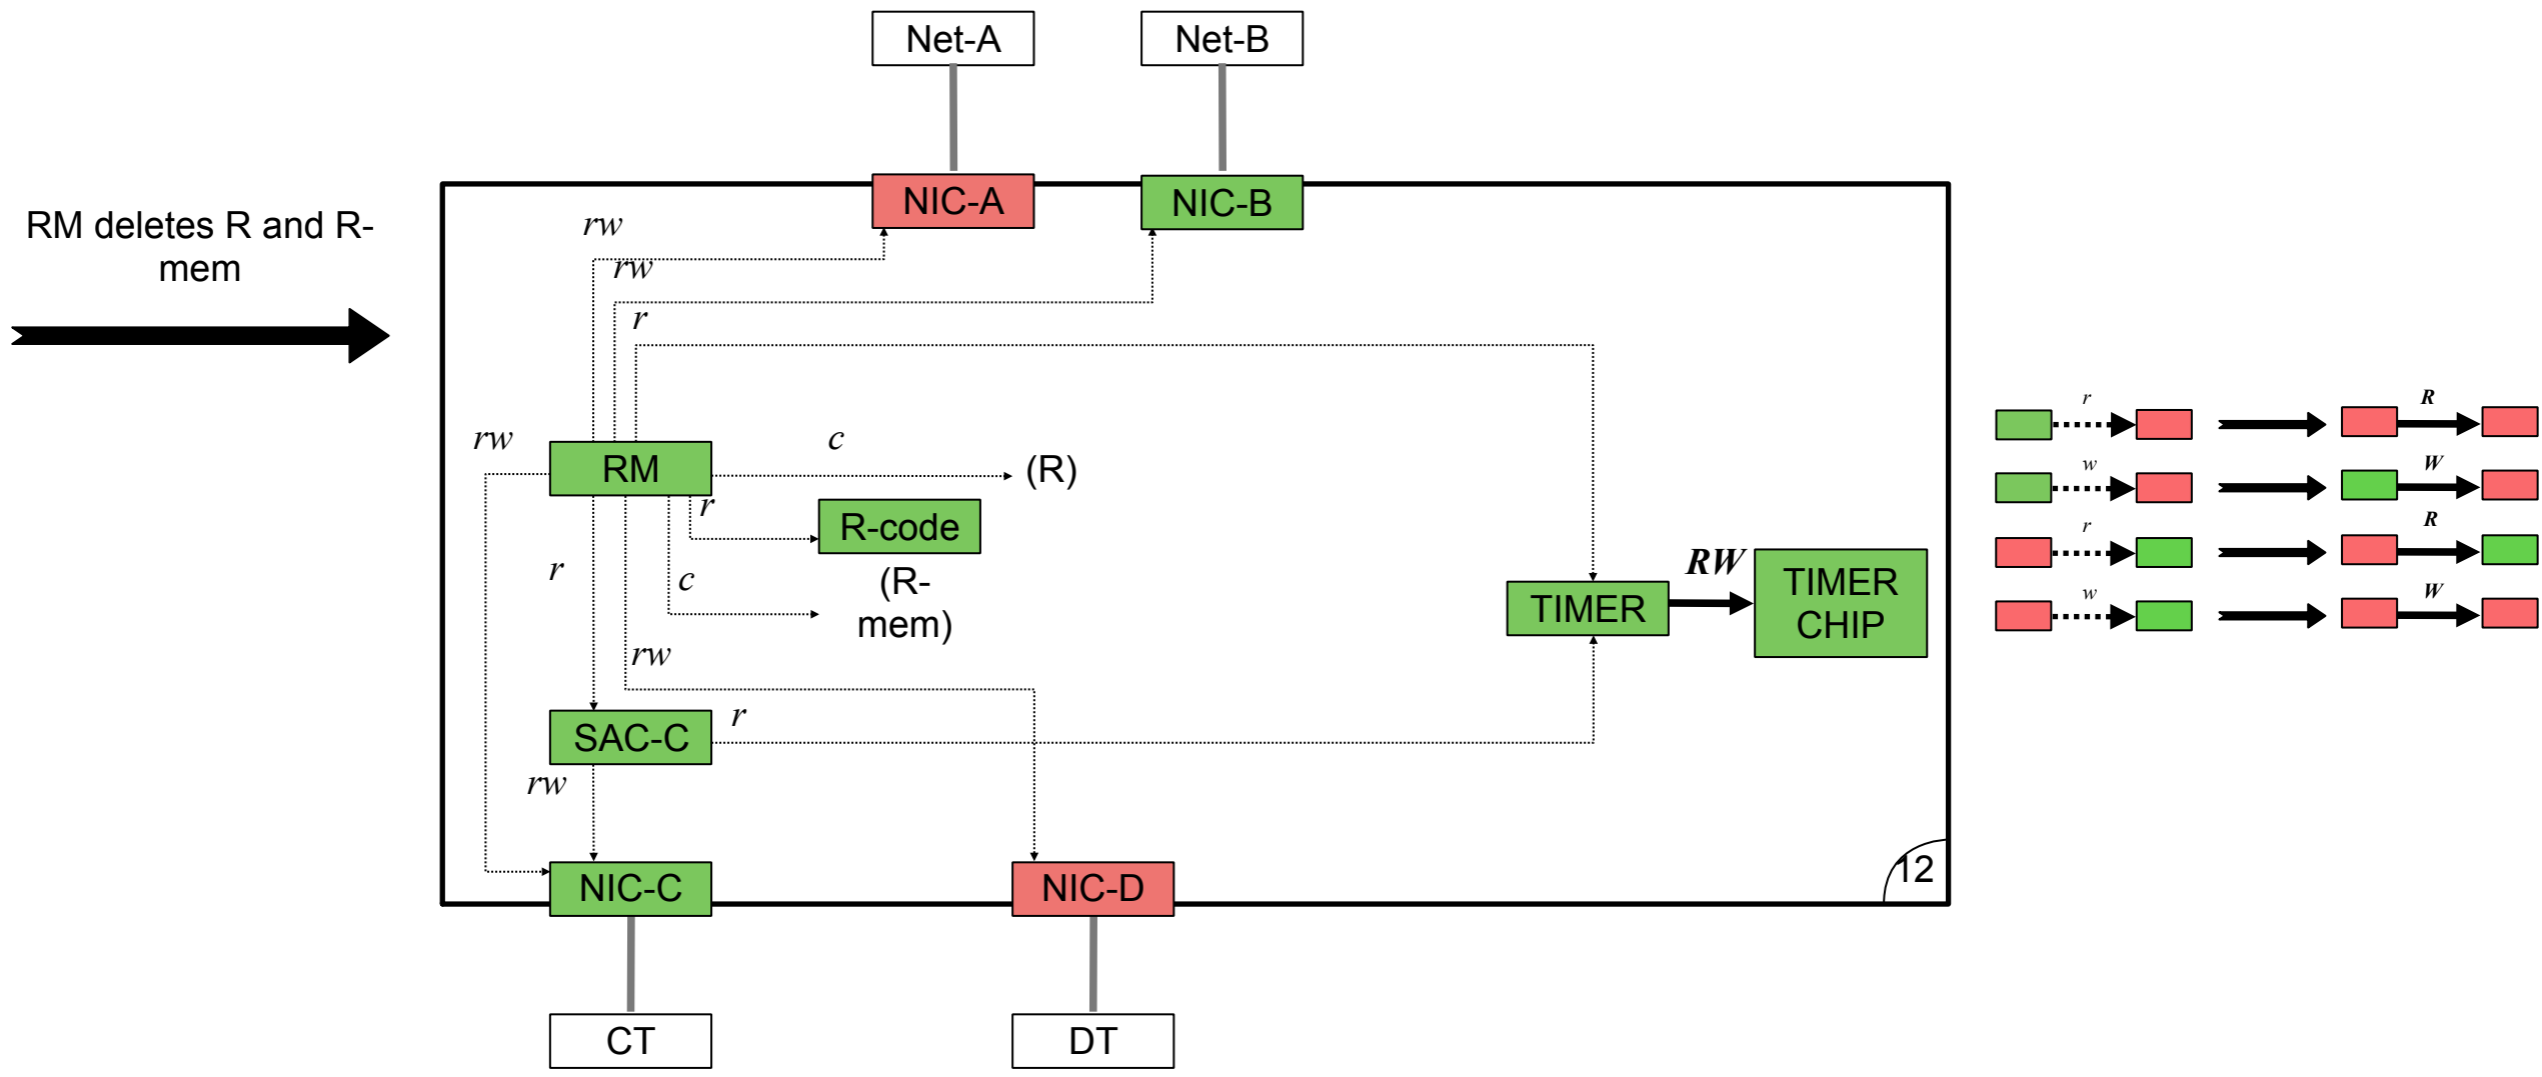


RM revokes caps of R-mem  
 (using create cap to R-mem)






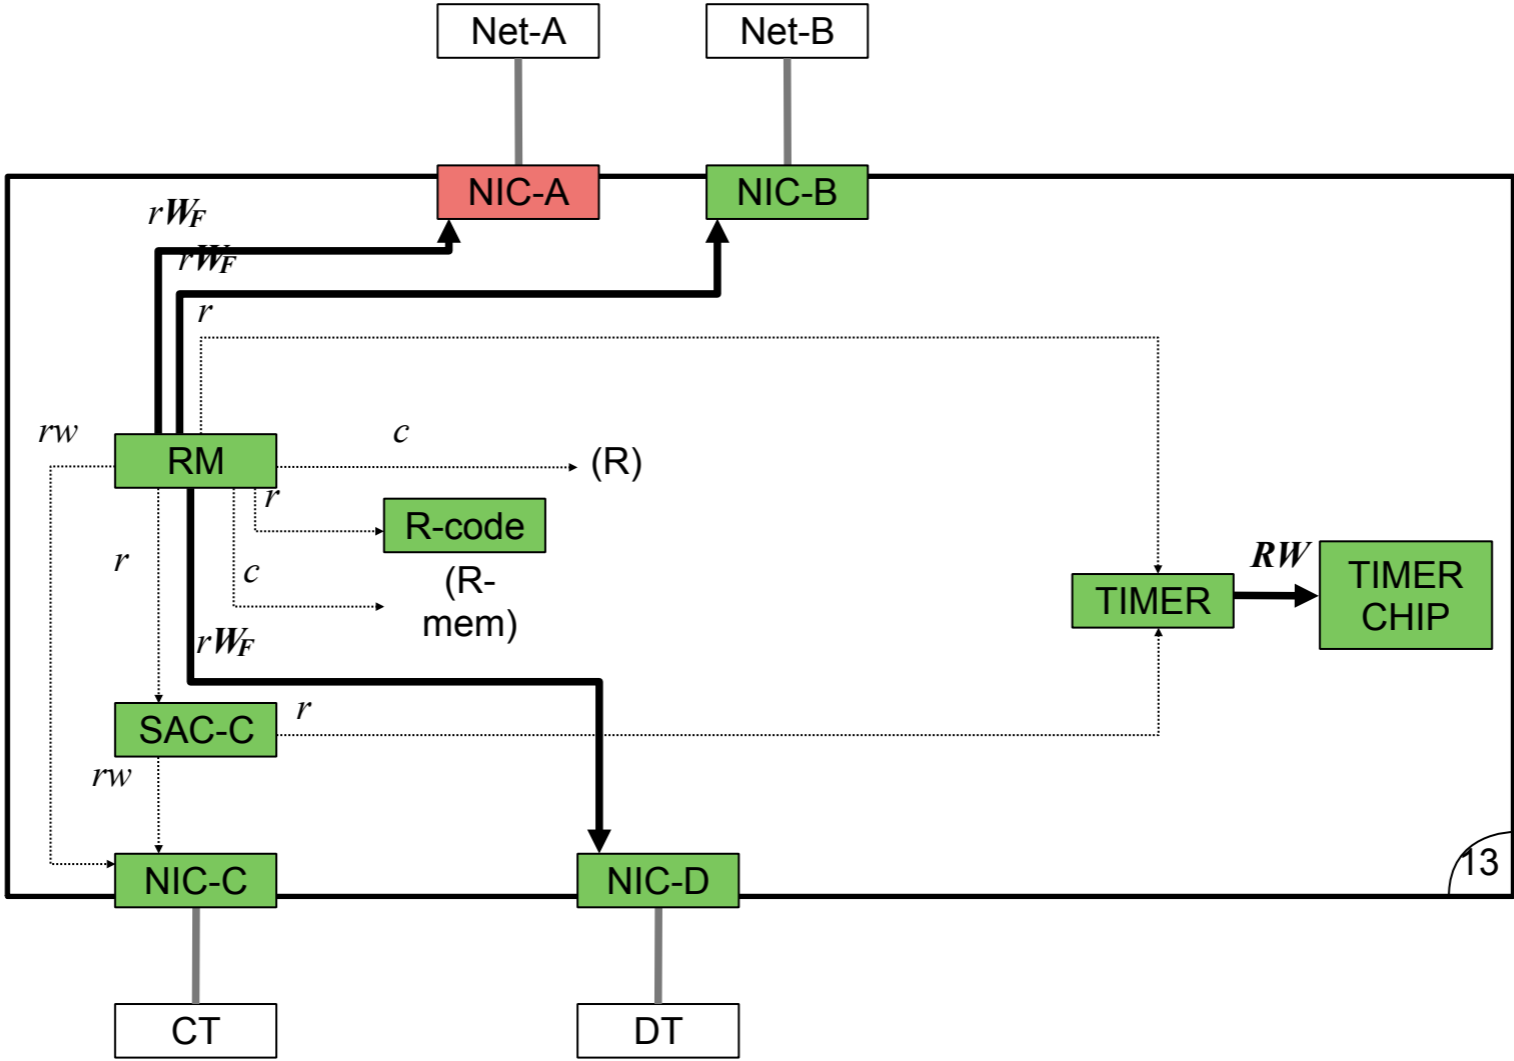
# Life Cycle



# Life Cycle


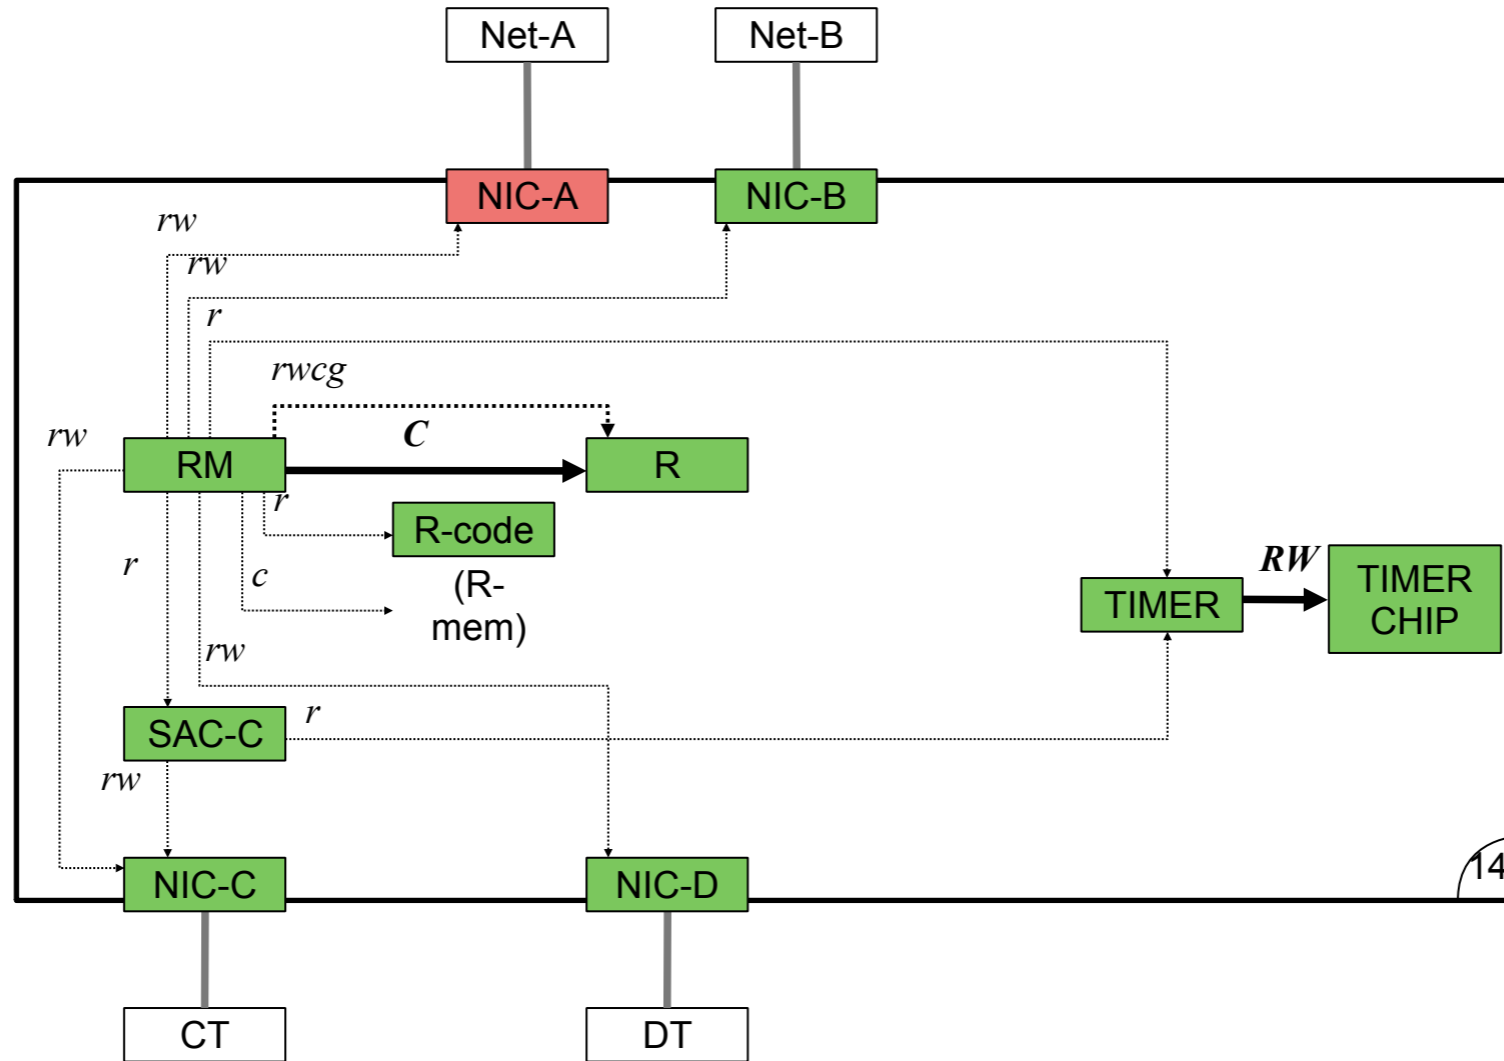


RM flushes  
NIC-A, NIC-B and  
NIC-D



# Life Cycle


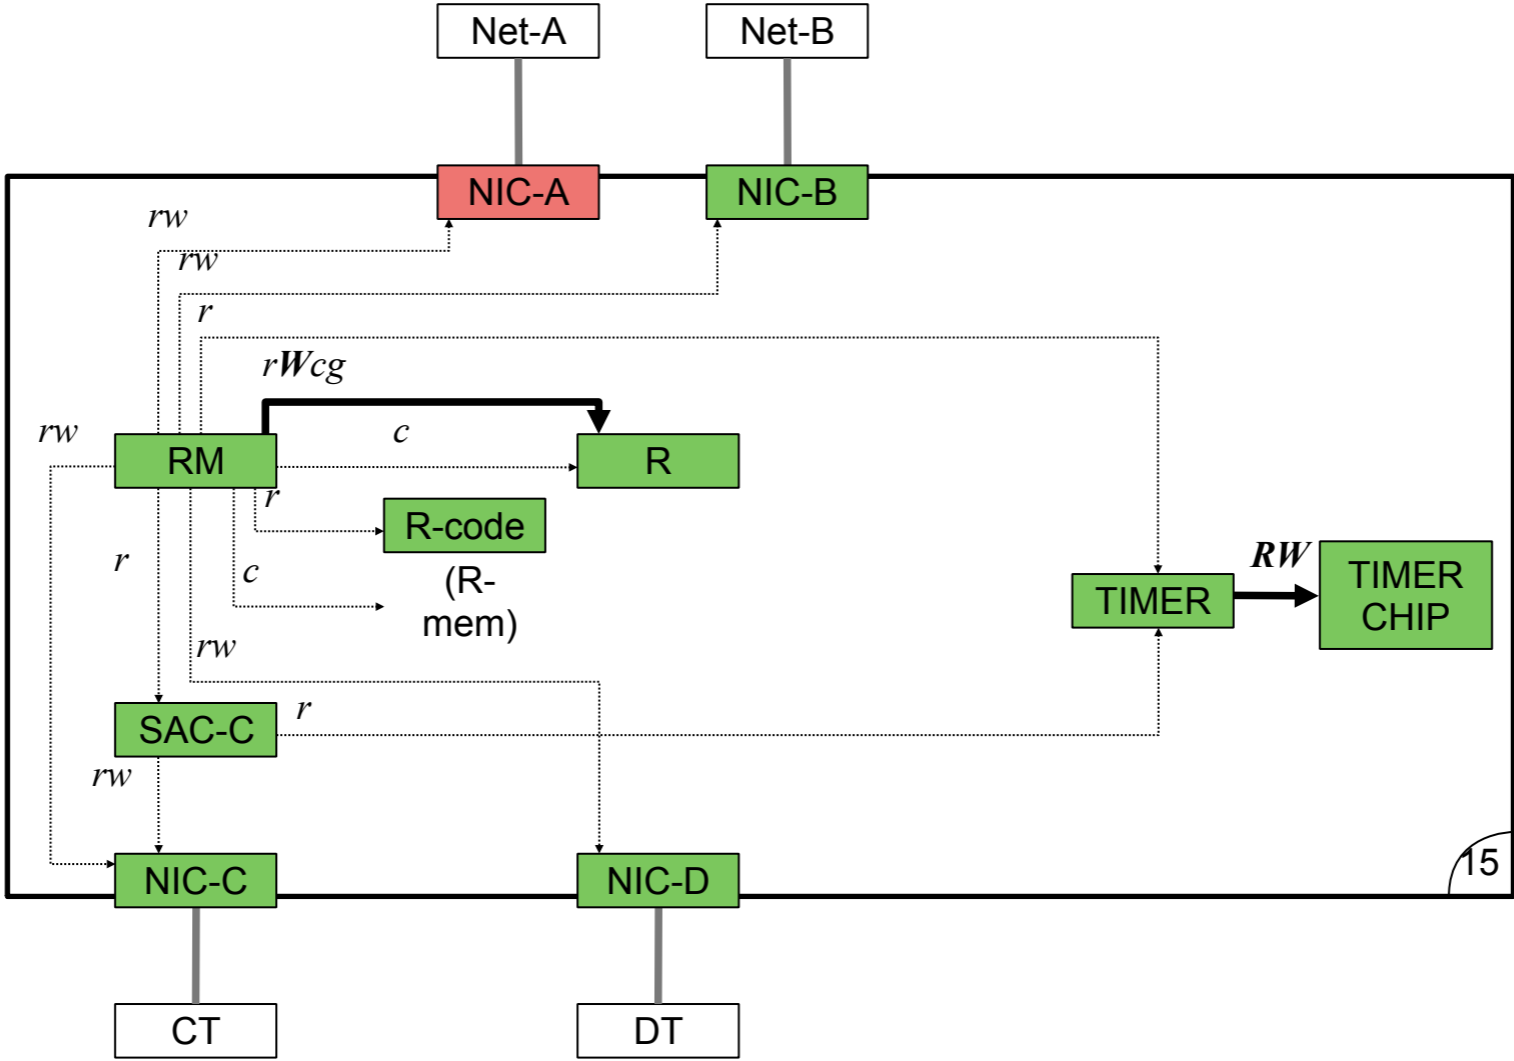
RM creates a new Router instance R  
 (and gets full rights to the newly created object)



# Life Cycle




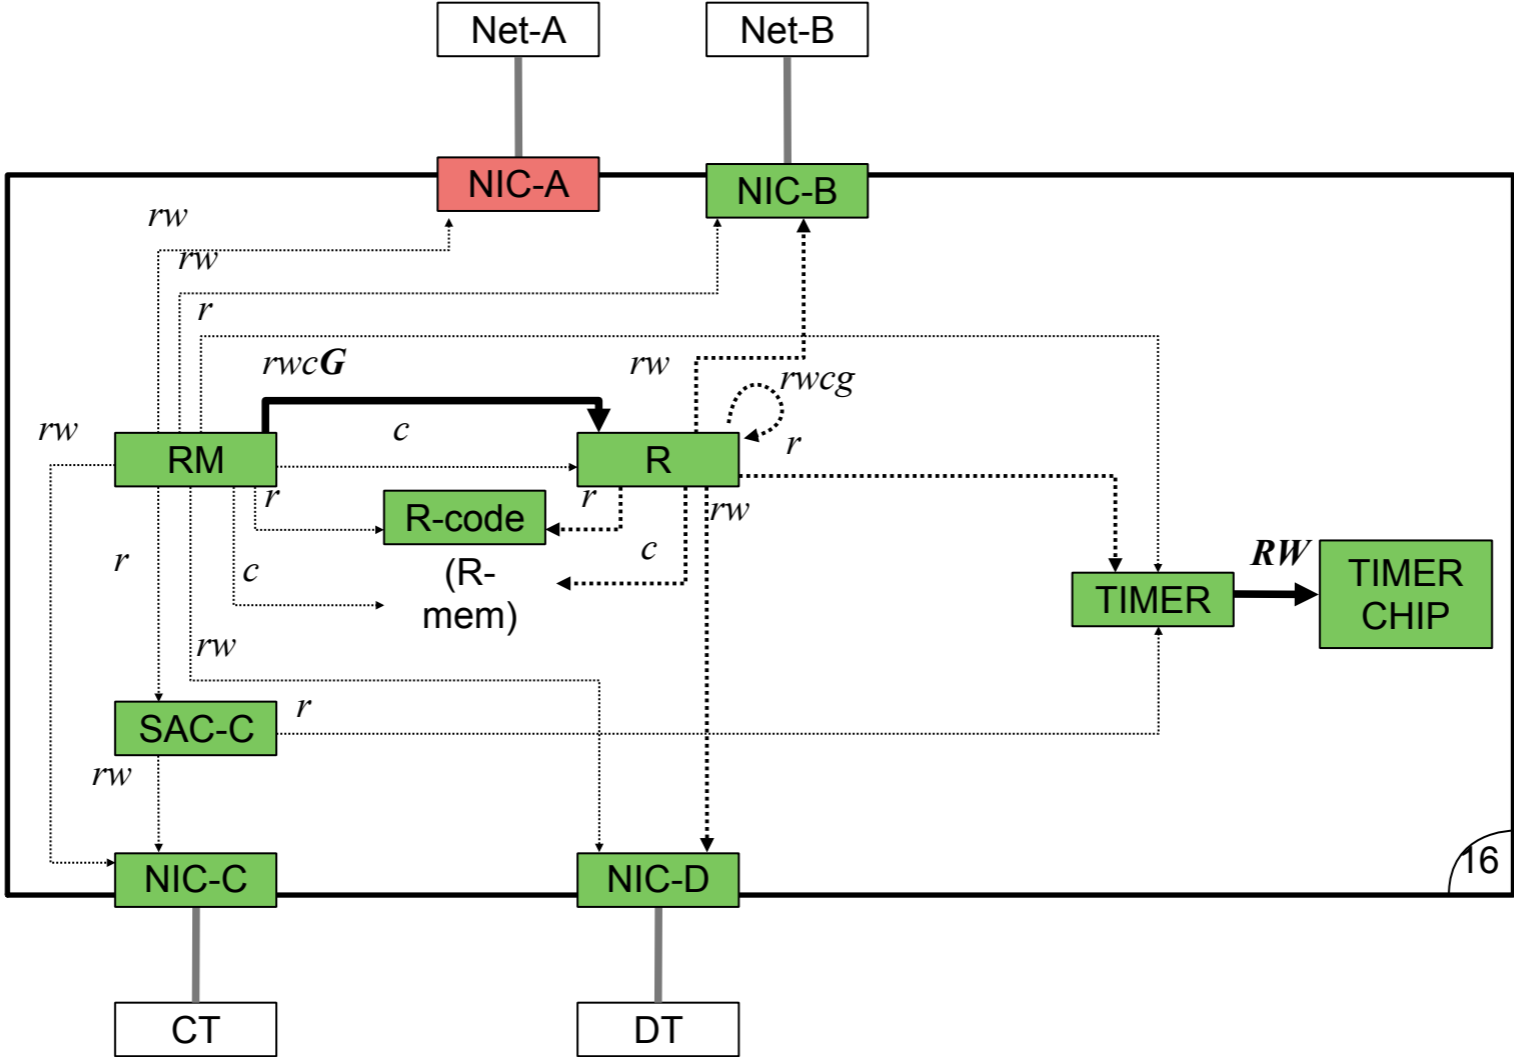
RM grants to R its rights to NIC-B, NIC-D, R-mem, R-code, TIMER and itself



# Life Cycle



RM grants to R its rights to NIC-B, NIC-D, R-mem, R-code, TIMER and itself


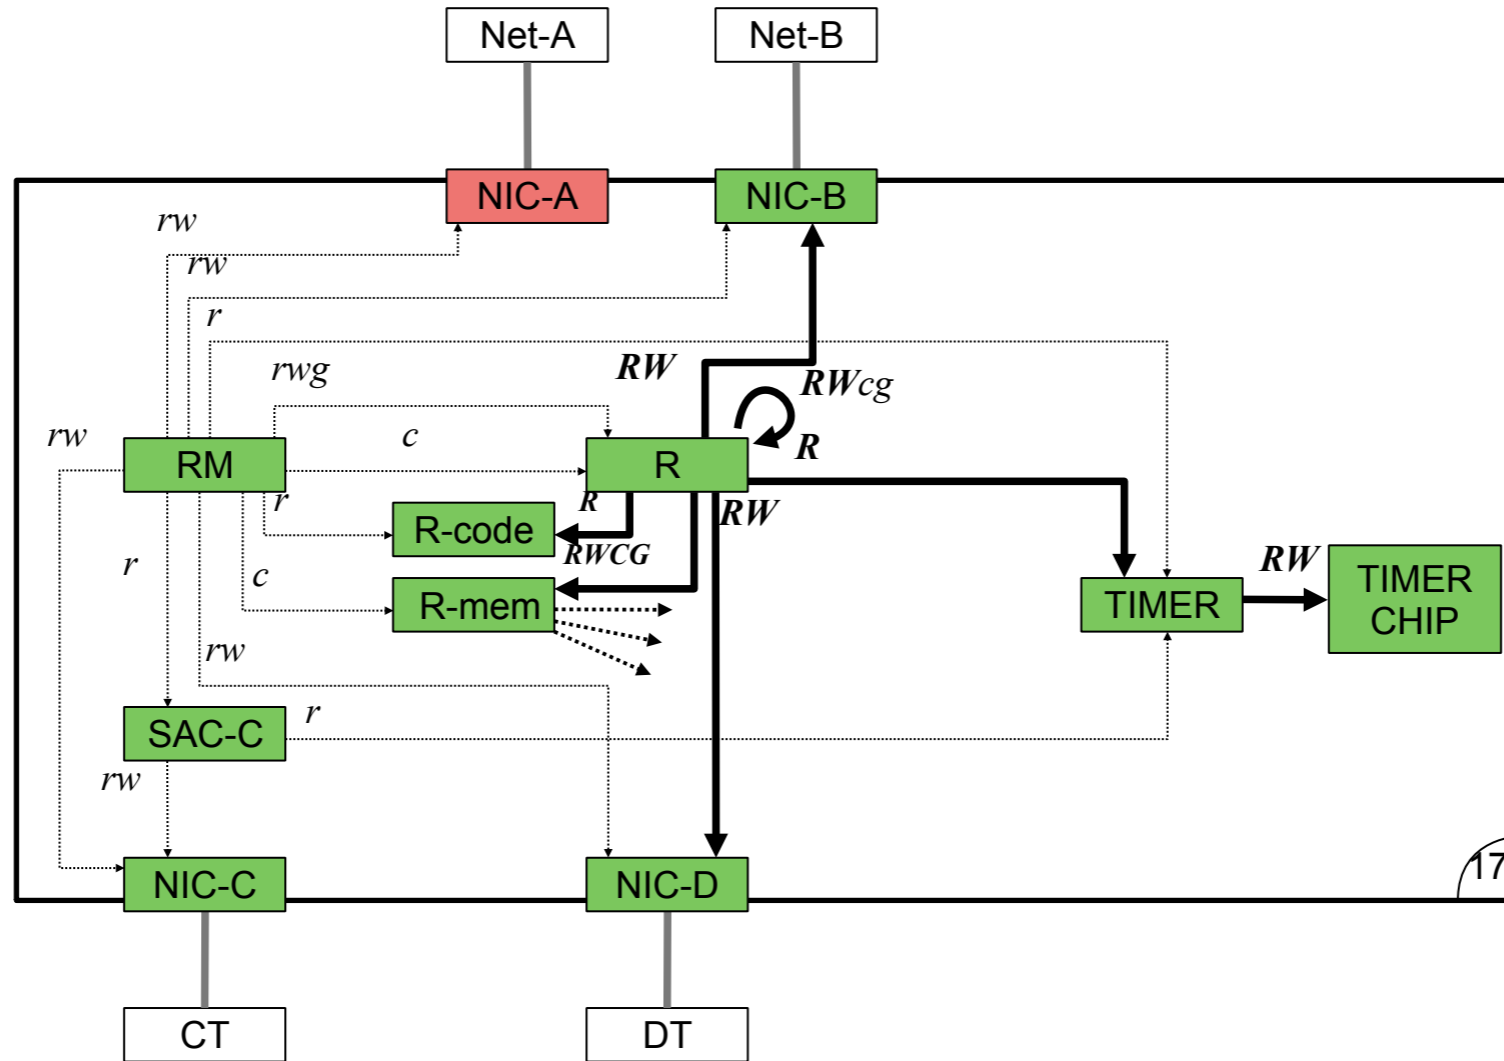


# Life Cycle


DT starts to communicate with Net-B, through R.

**→**

This may imply the creation of new objects using R-mem and granting some caps to them.



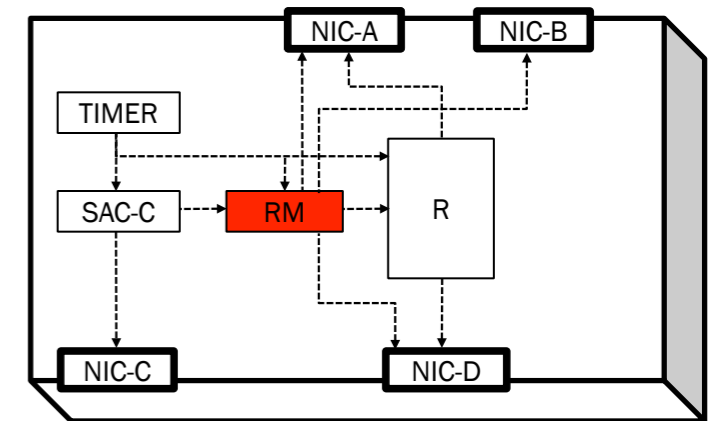
# So far







# So far

- Can build systems with
  - large untrusted components
  - plus few small, trusted components
  - trusted = needs behaviour spec

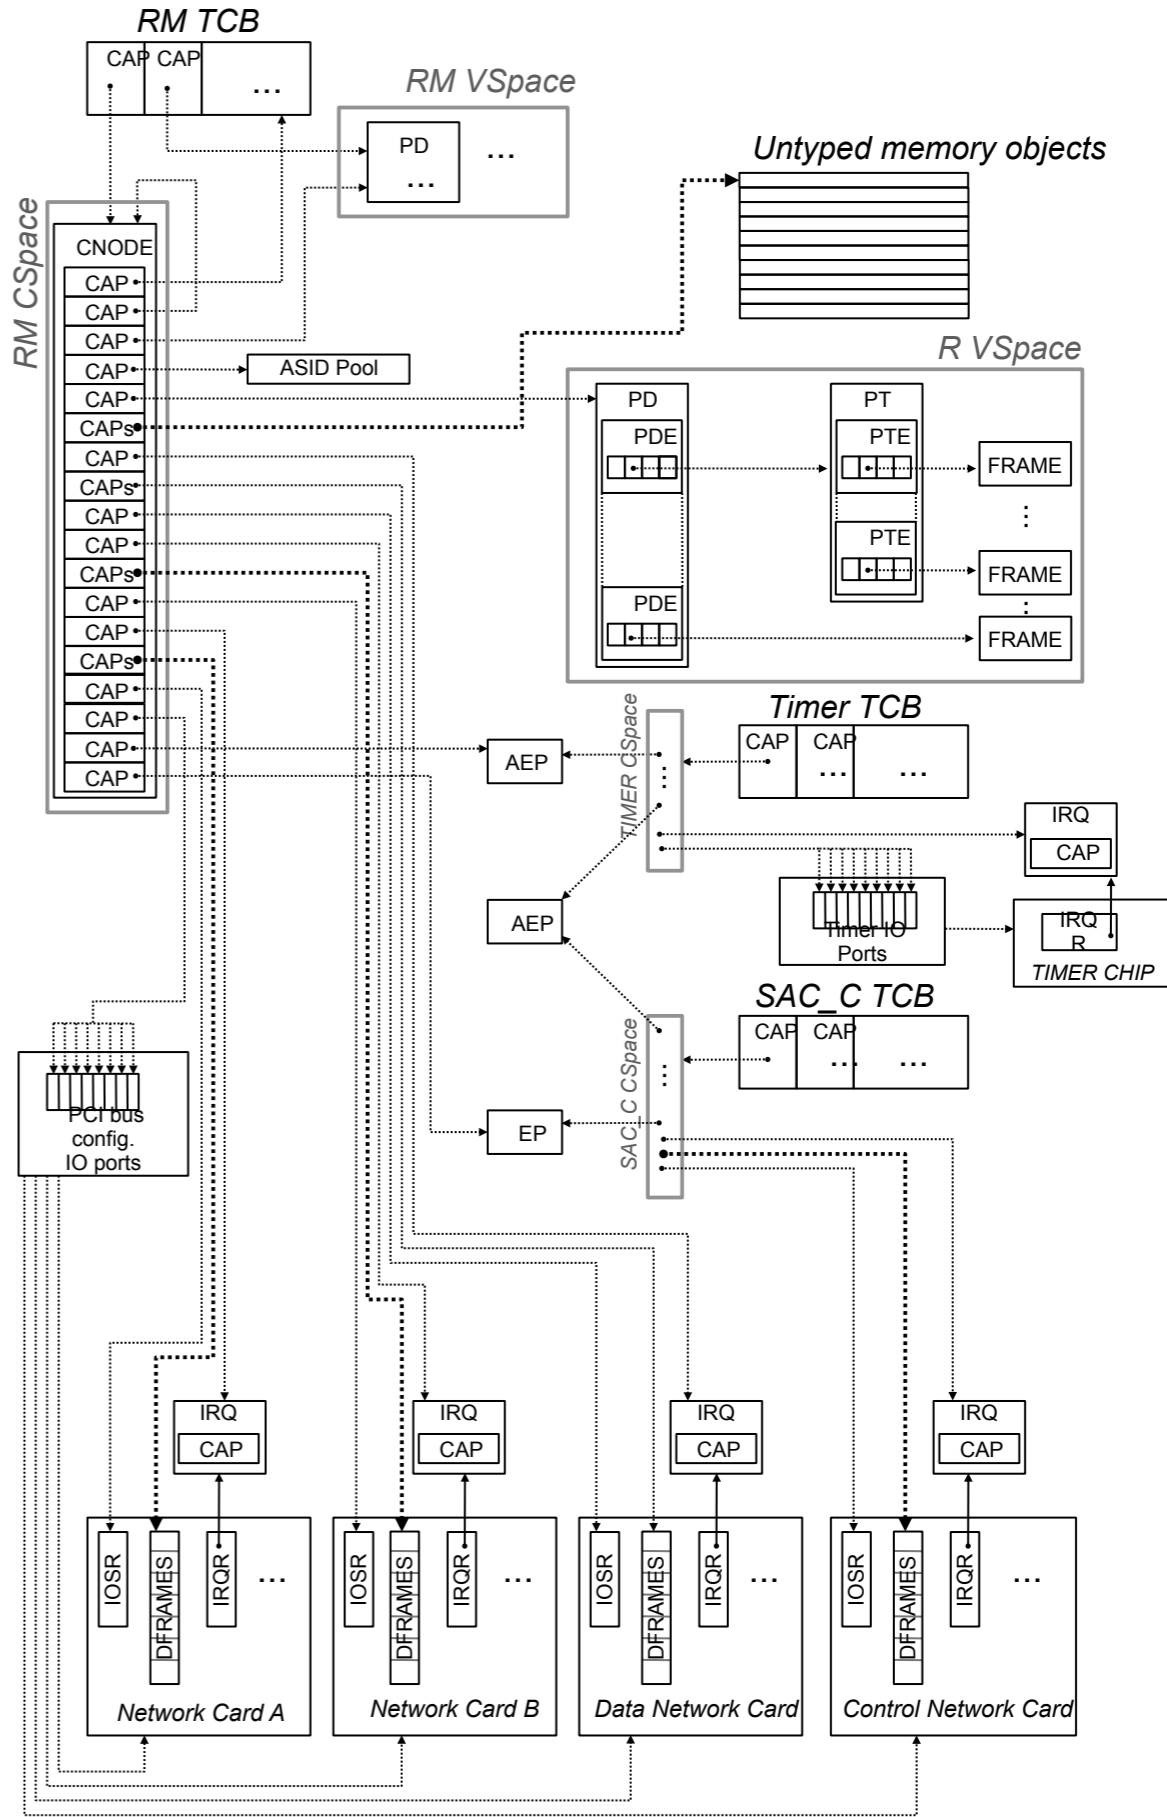


# So far

- Can build systems with
  - large untrusted components
  - plus few small, trusted components
  - trusted = needs behaviour spec
- Use take-grant to model security
  - can simulate system
  - modelling already finds bugs
  - high-level proof in Isabelle/HOL or SPIN
  - includes behaviour of trusted component




# Future



IOSR = IOSpace Root  
 Pointer  
 IRQR = IRQ register  
 reference  
 DFRAMES = Device Frames

# Future

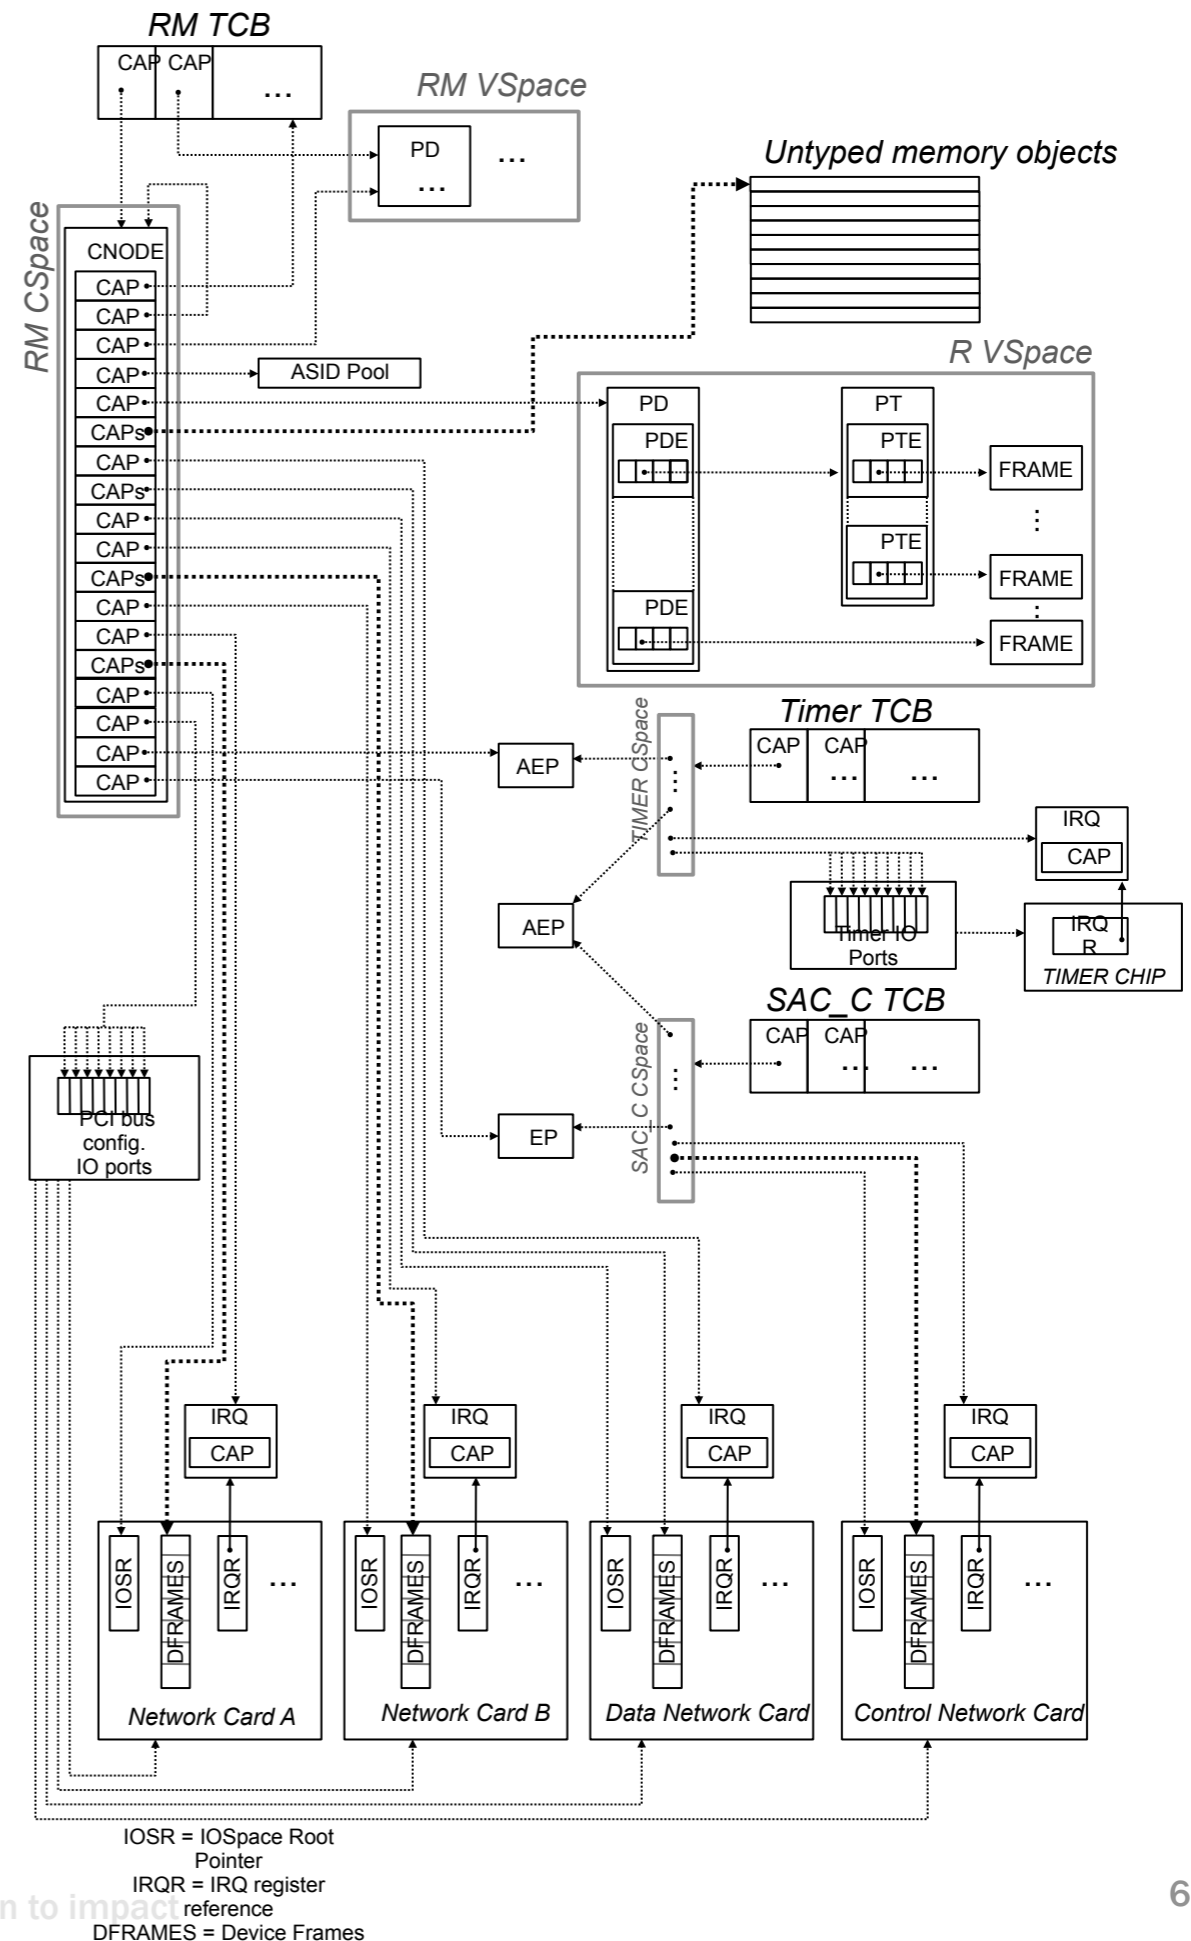
- Need to verify low-level design




IOSR = IOSpace Root Pointer  
 IRQR = IRQ register reference  
 DFRAMES = Device Frames

# Future

- Need to verify low-level design
- Building tool-chain for:
  - describing cap layout (capDL)
  - generating booter
  - generating booter proof
  - abstraction to take-grant

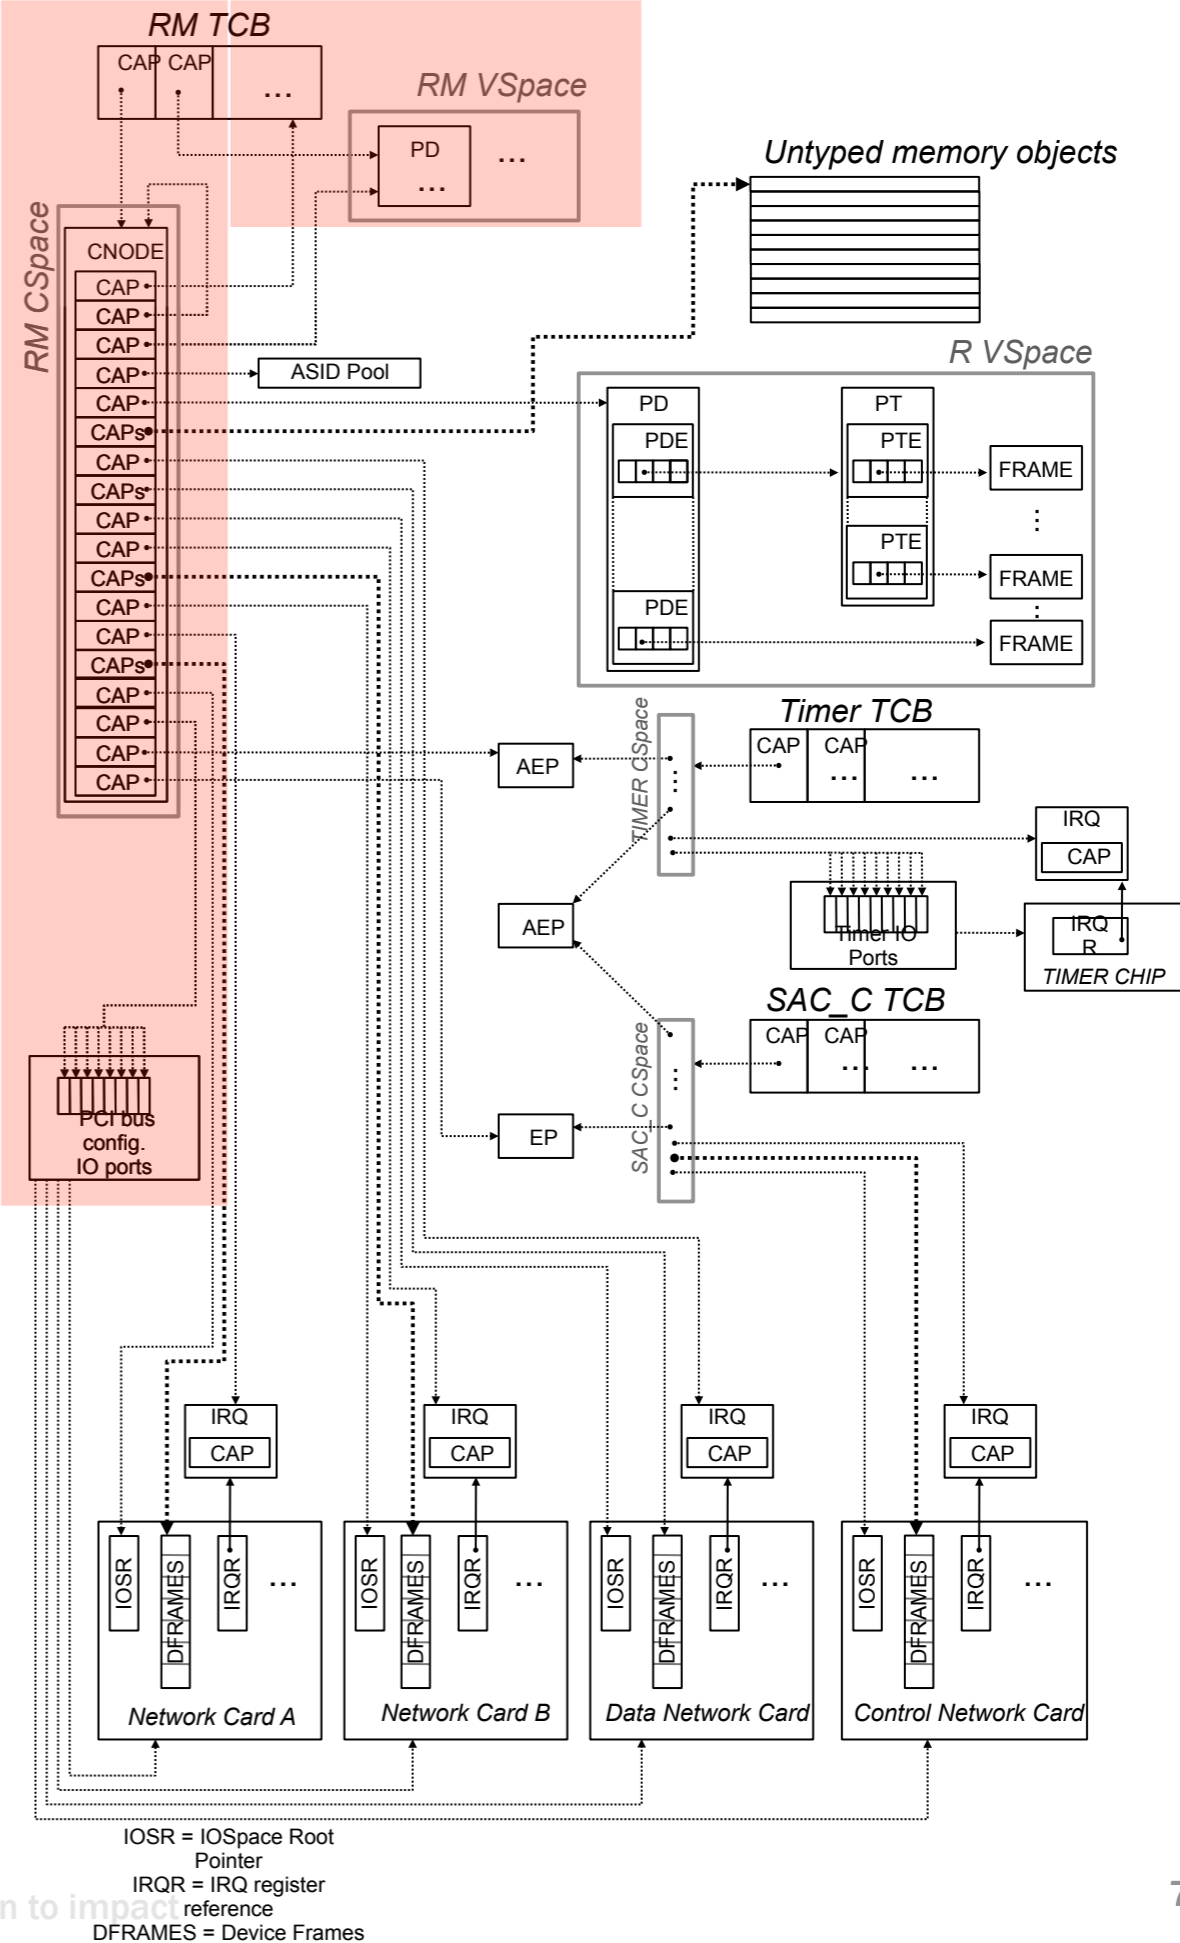


# More Future



# More Future

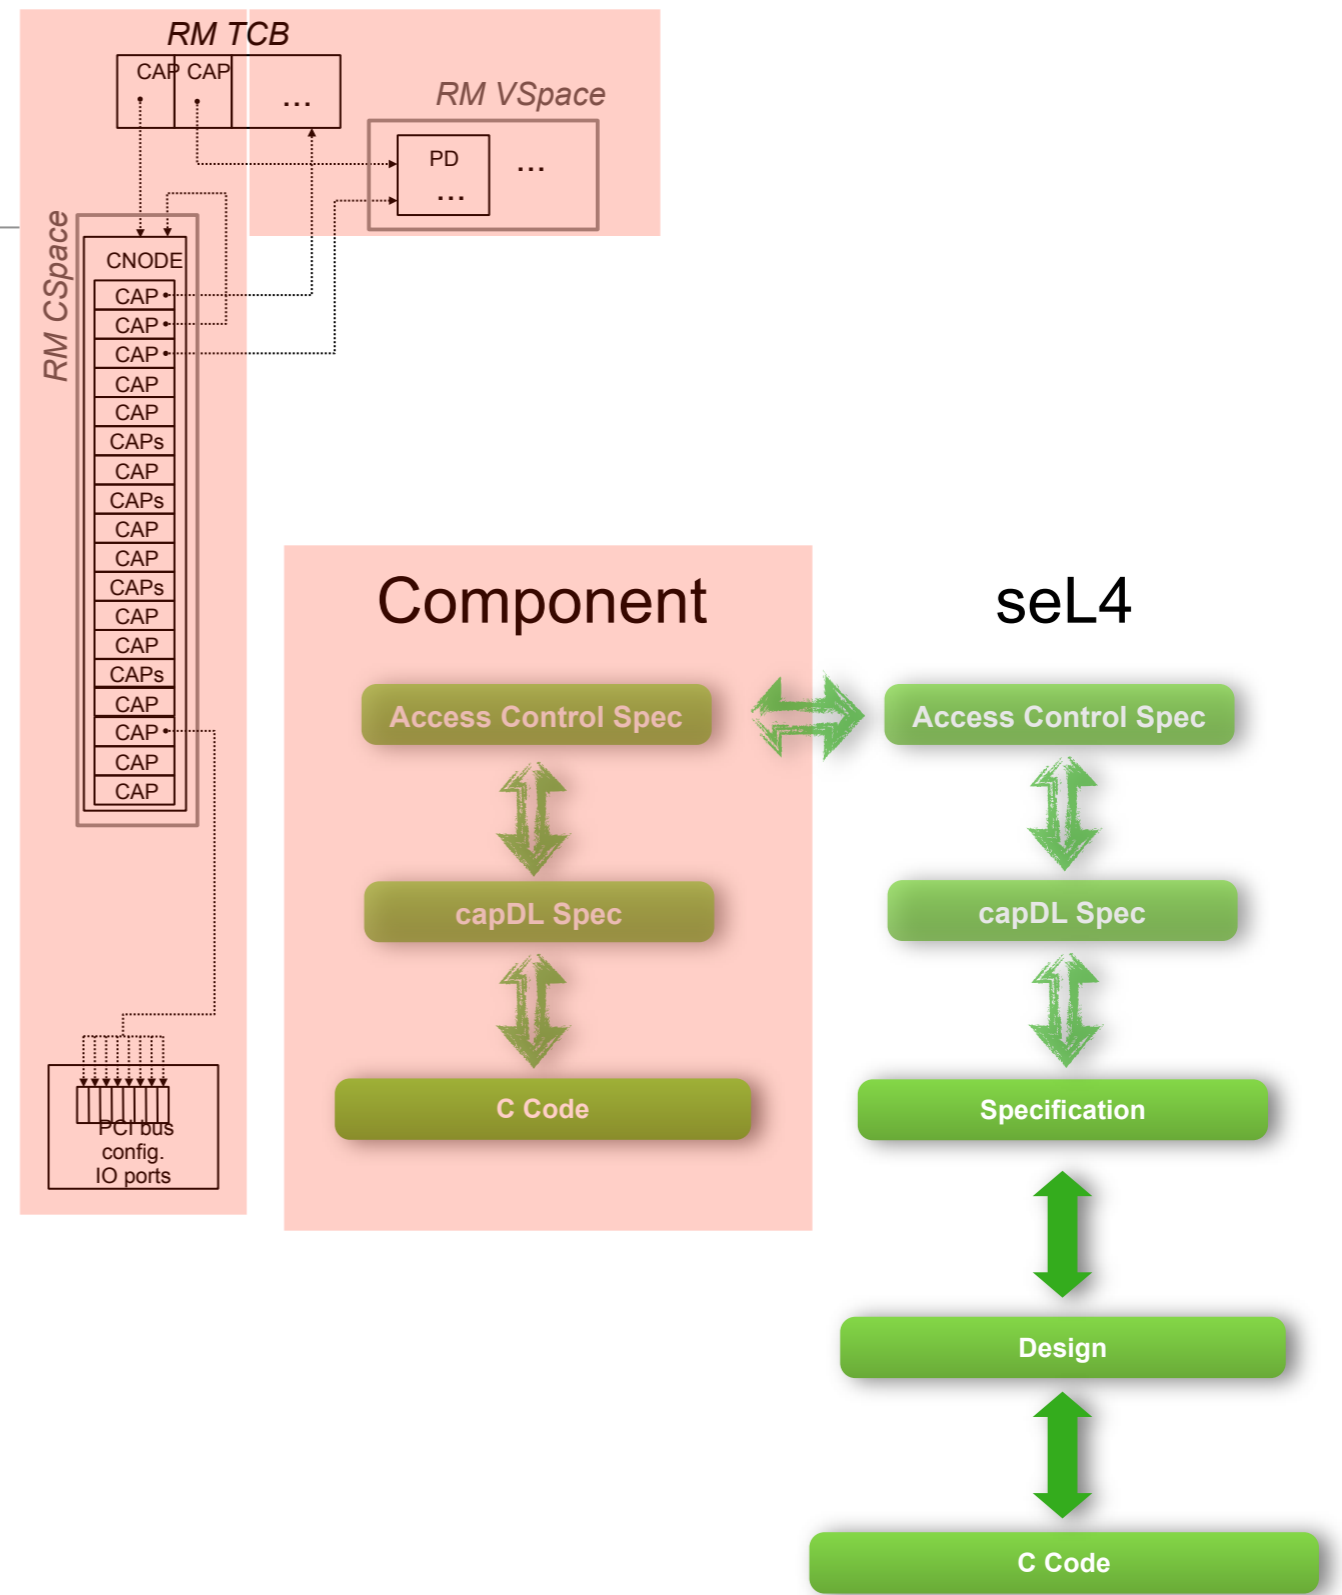
- Verify Trusted Component





# More Future

- **Verify Trusted Component**
- **Refine to C:**
  - interface with kernel
  - use most abstract level possible
  - make sure sec property preserved by refinement



# Summary



# Summary



# Summary

Formal proof all the way from spec to C.

- **200kloc** handwritten, machine-checked proof
- **~460** bugs (160 in C)
- Verification on **code, design, and spec**
- Systems with **trusted components**
- **The future: formal proof for large systems down to code**



Formal Code Verification up to 10kloc:

It works.  
It's feasible.  
It's fun.







**Thank You**



# Thank You

Google