Progress Towards a Trustworthy Systems Platform

Gernot Heiser
NICTA and University of New South Wales
Sydney, Australia

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

NICTA Funding and Supporting Members and Partners

Previously Reported: Formal Verification of seL4

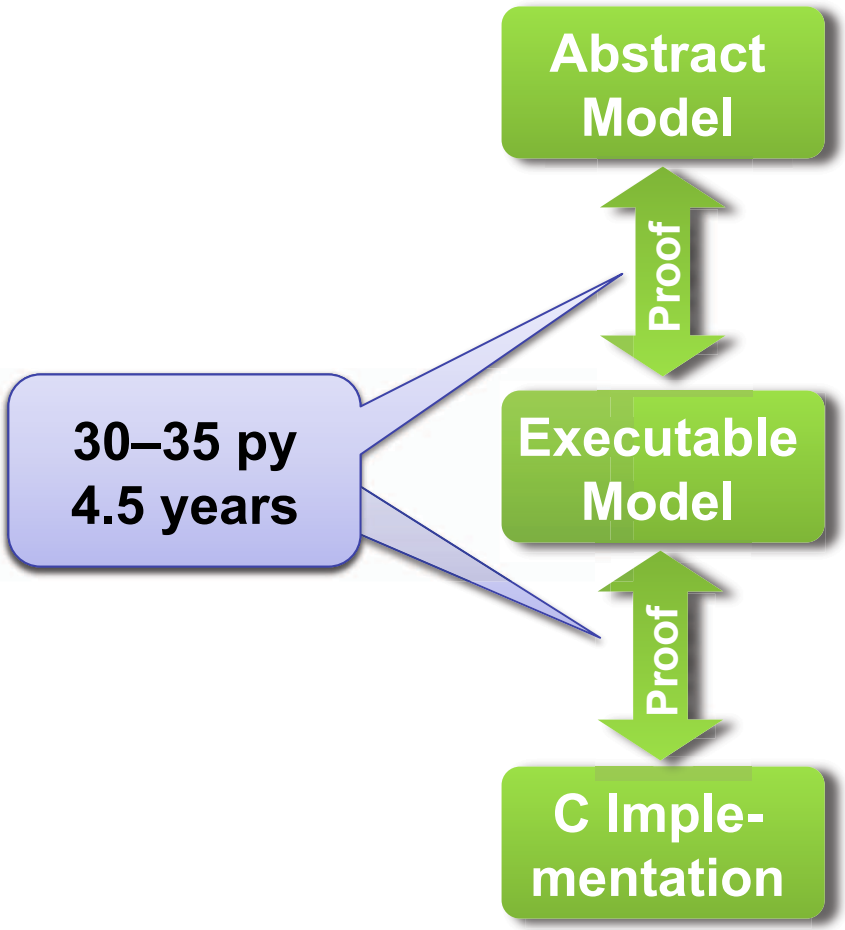

Result: a microkernel proved functionally correct

So what?

Goal: Trustworthiness/Dependability

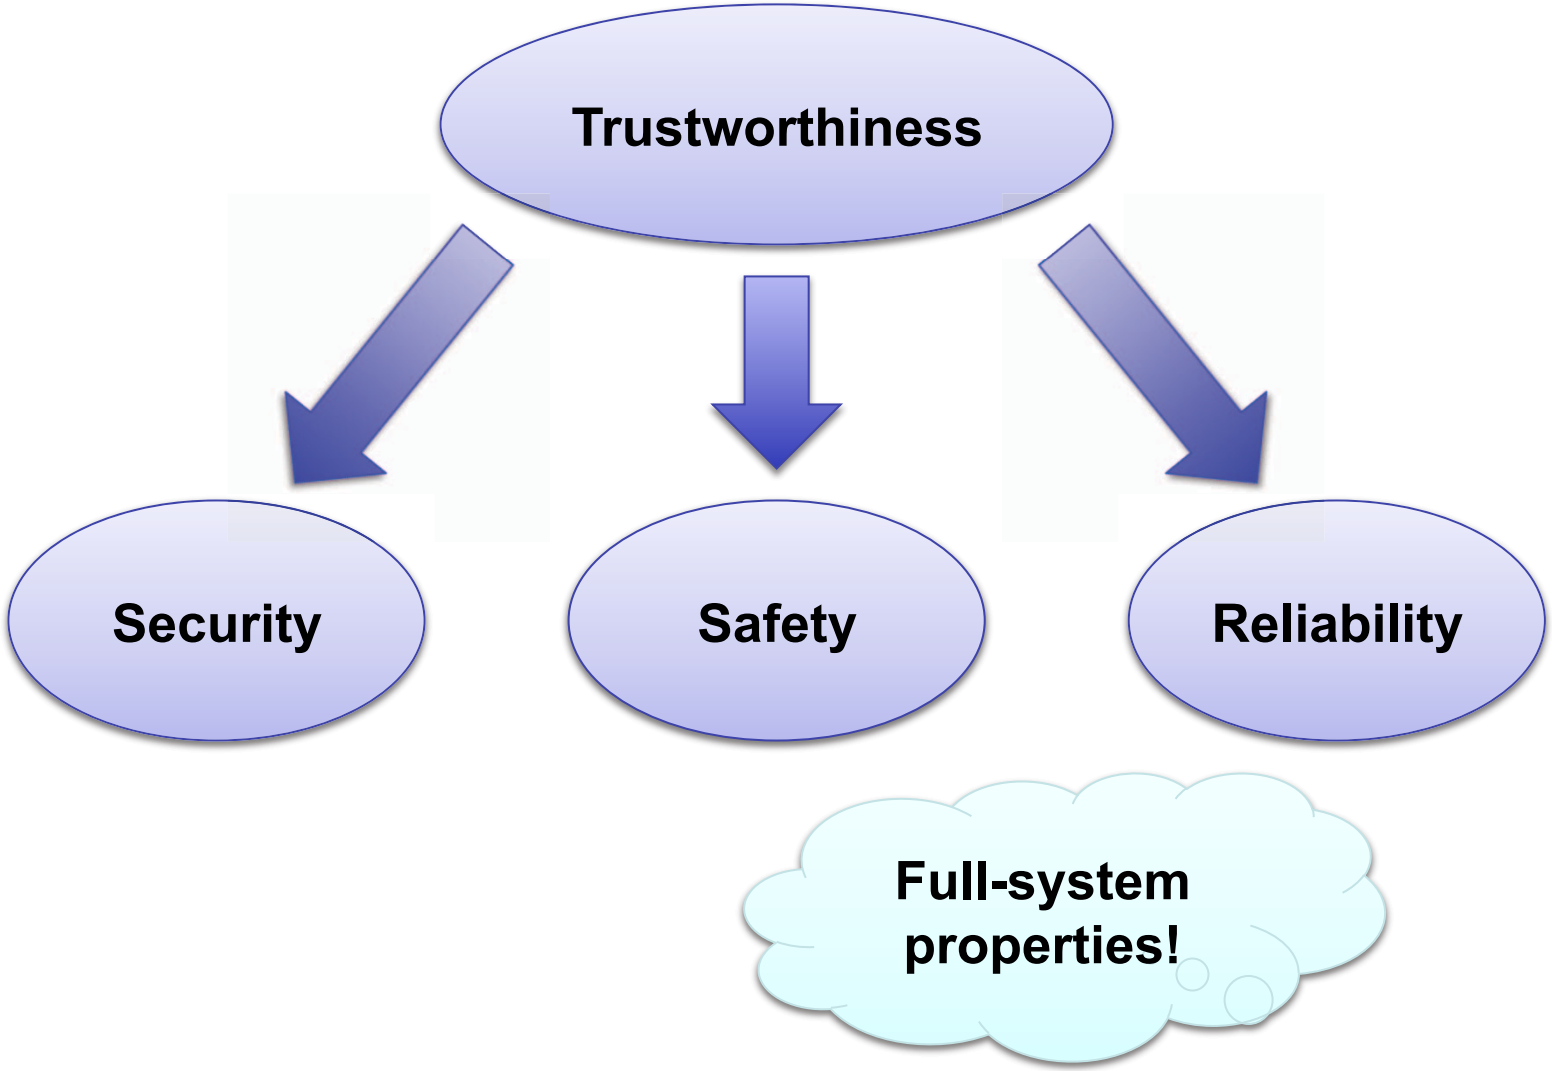

Real-World Trustworthiness

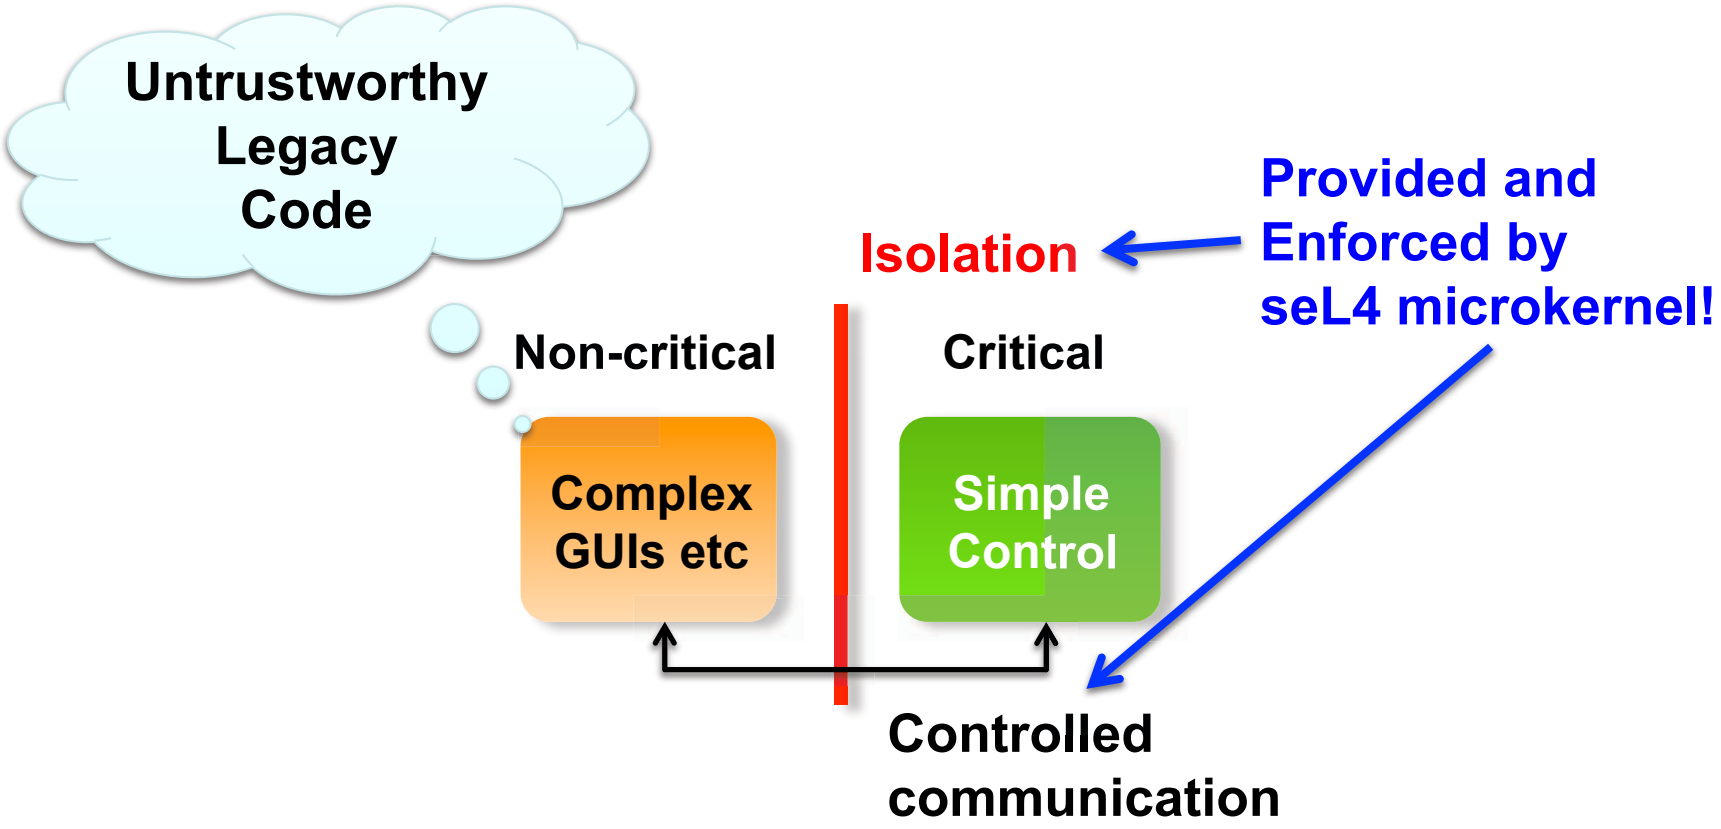

seL4 as Basis for Trustworthy Systems

NICTA

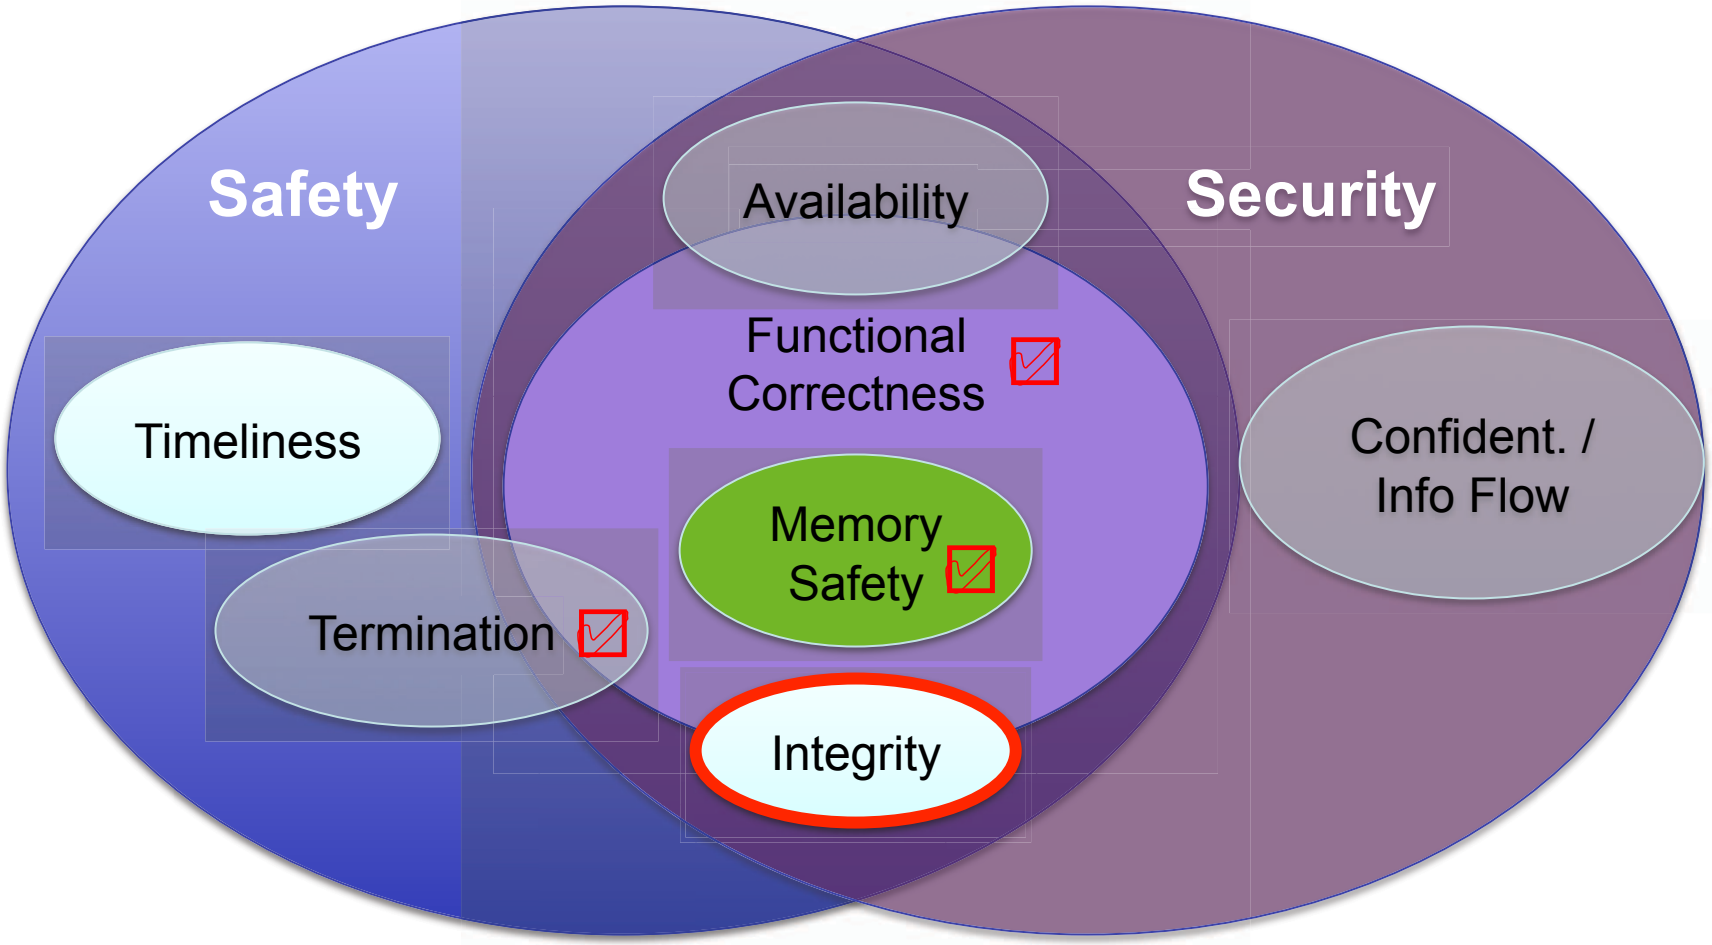

Integrity is about Write Accesses

NICTA

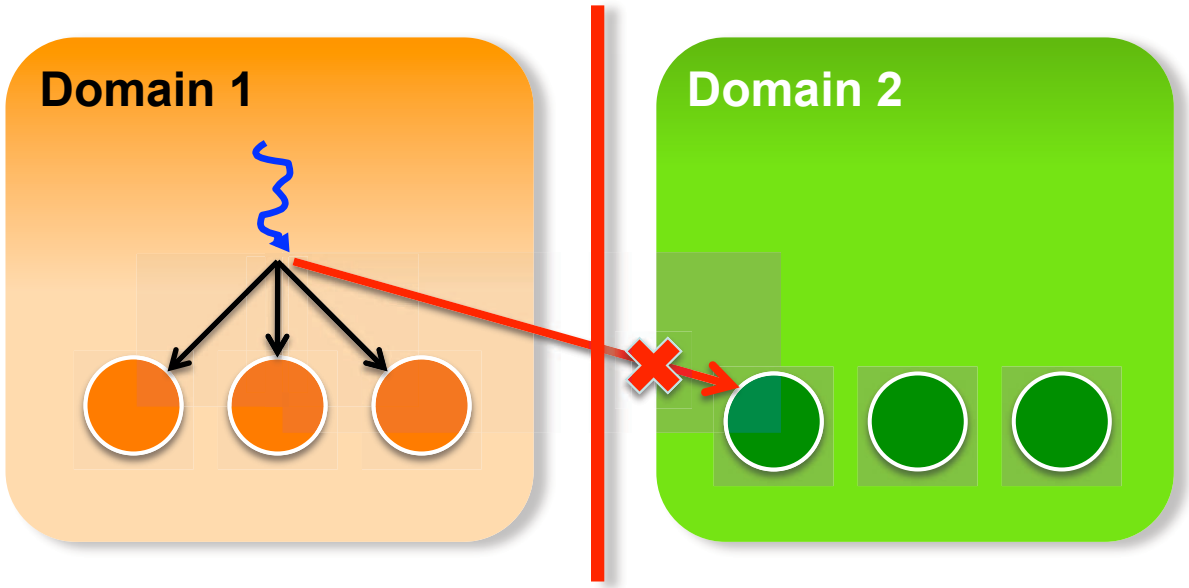

To prove:

- Domain-1 doesn't have write capabilities to Domain-2 objects
  ☐ no action of Domain-1 agents will modify Domain-2 state
- Specifically, kernel does not modify on Domain-1's behalf!

Simplified by seL4 Resource Management

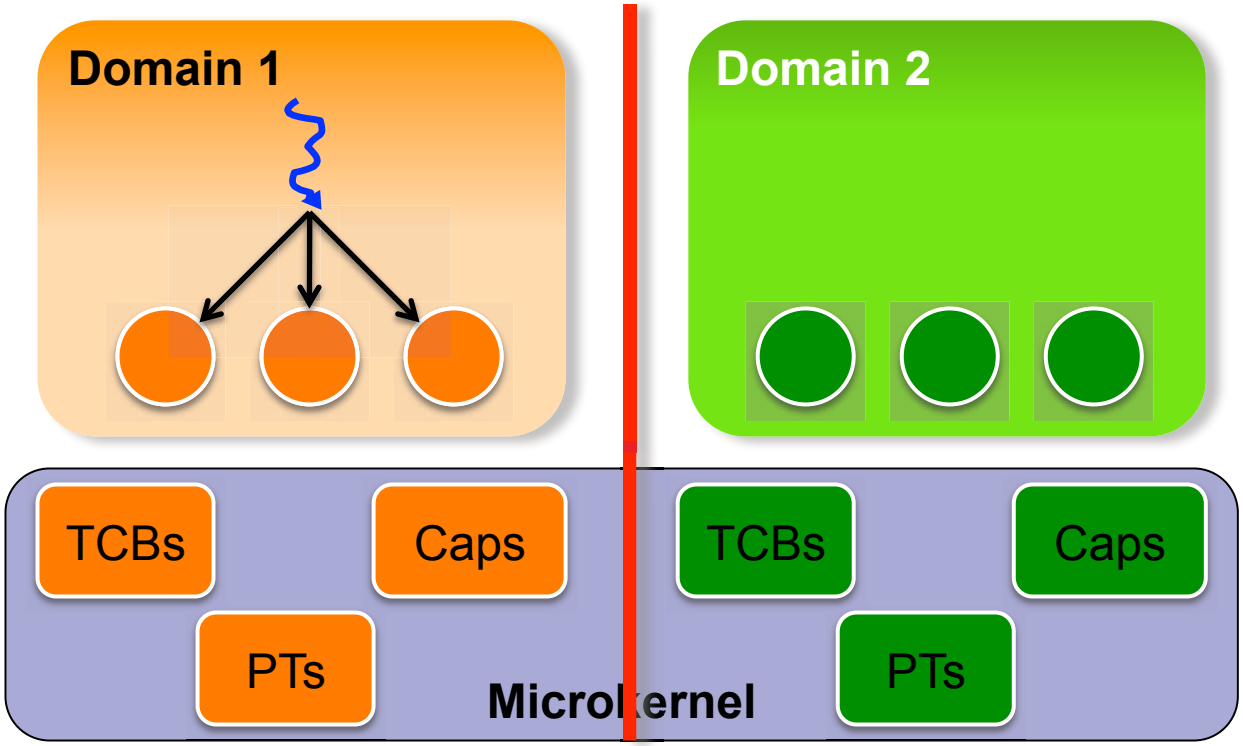

- Kernel data structures allocated by user
  - Protected by caps just as user data!
- Must show that no object can be modified without a write cap
  - Done last year [ITP'11], seL4 is first OS kernel with such a proof

seL4 as Basis for Trustworthy Systems

NICTA

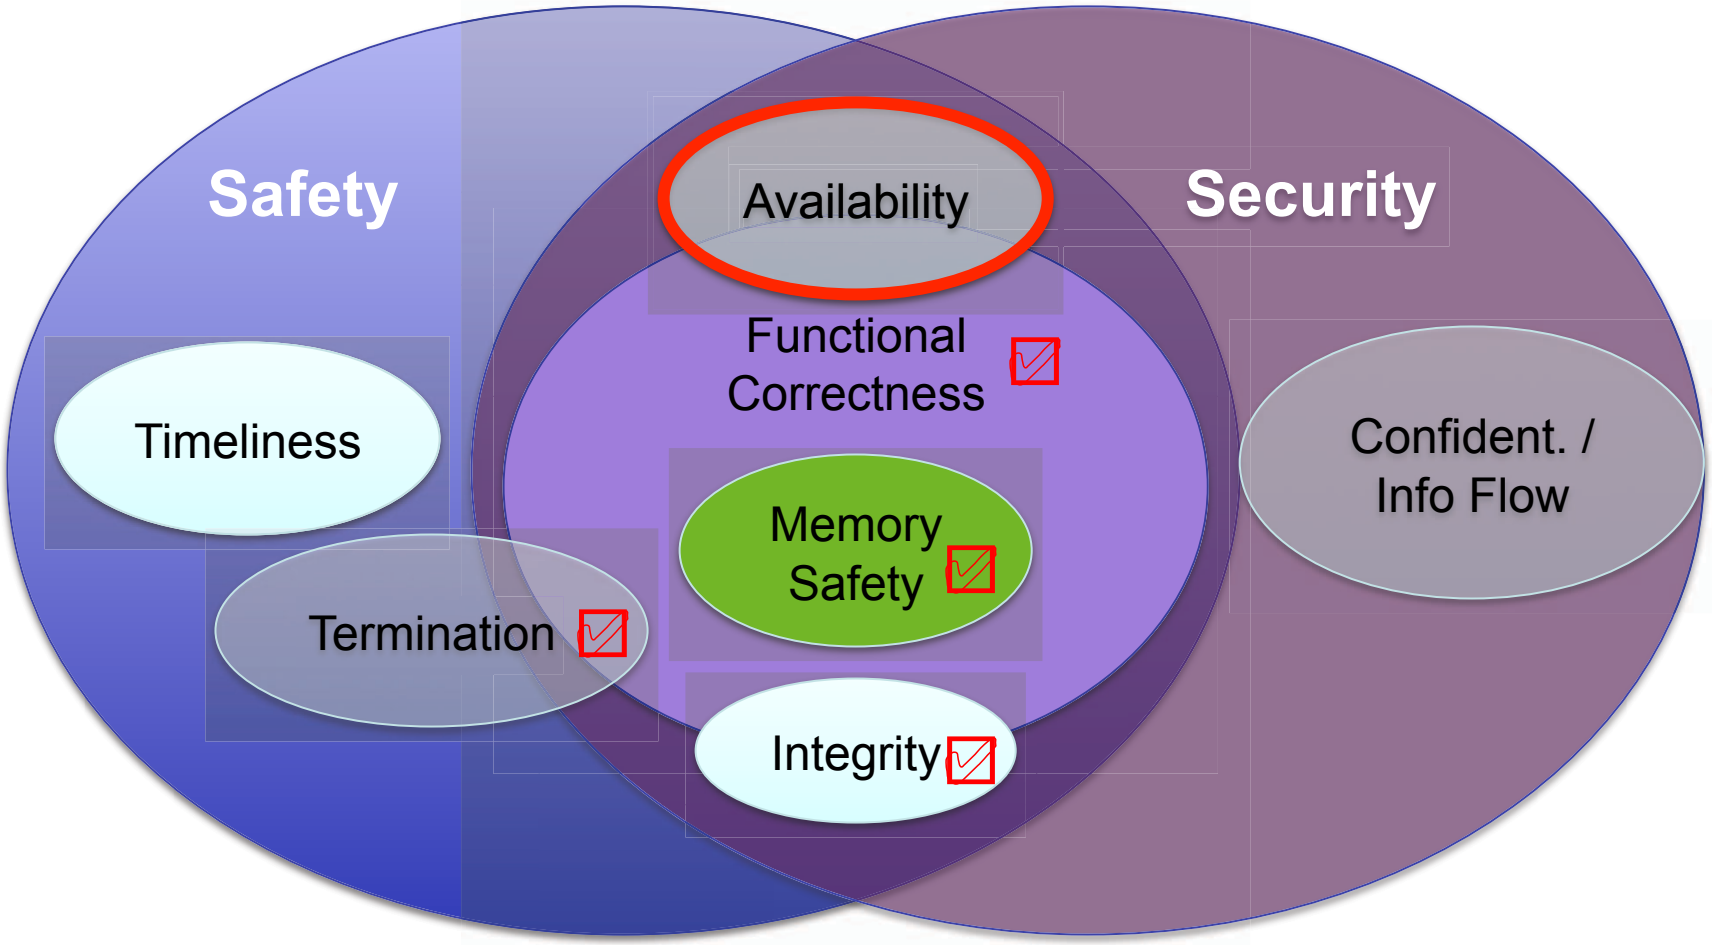

Availability is Trivially Ensured at Kernel Level

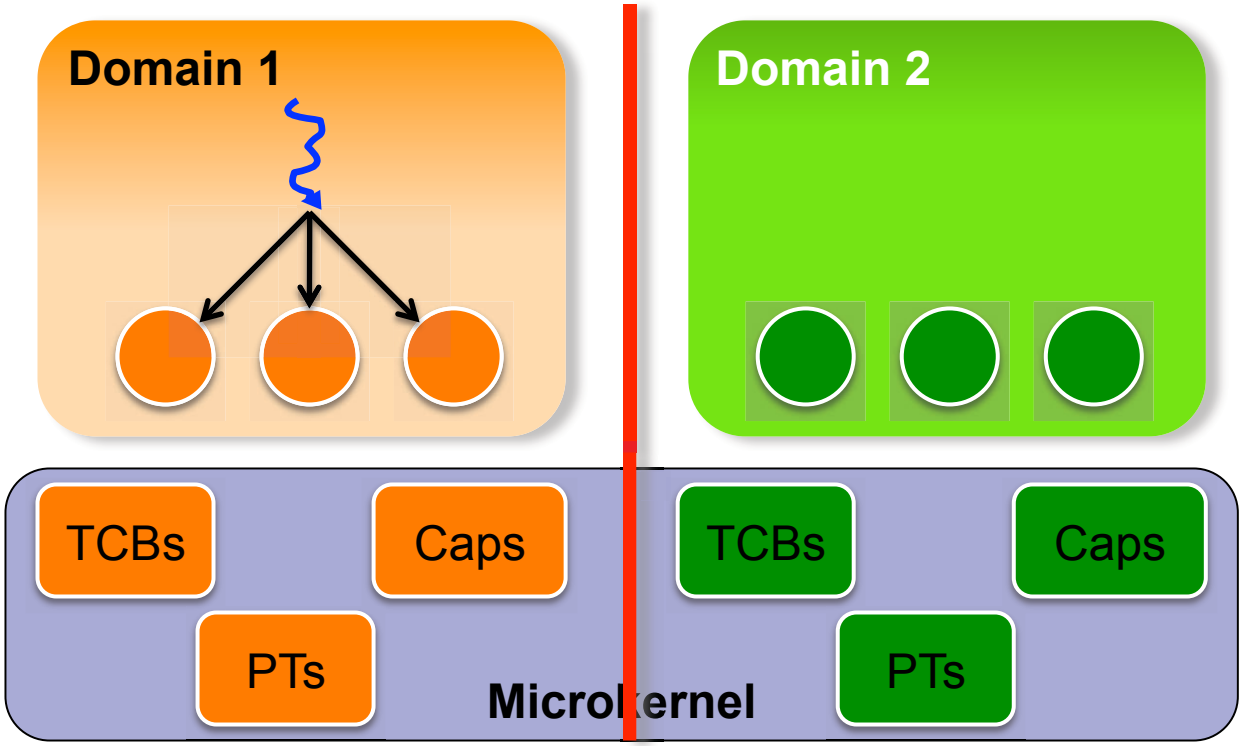

- Strict separation of kernel resources
  ☐ agent cannot deny access to another domain's resources
- Managing resource availability is a user-level issue

seL4 as Basis for Trustworthy Systems

NICTA

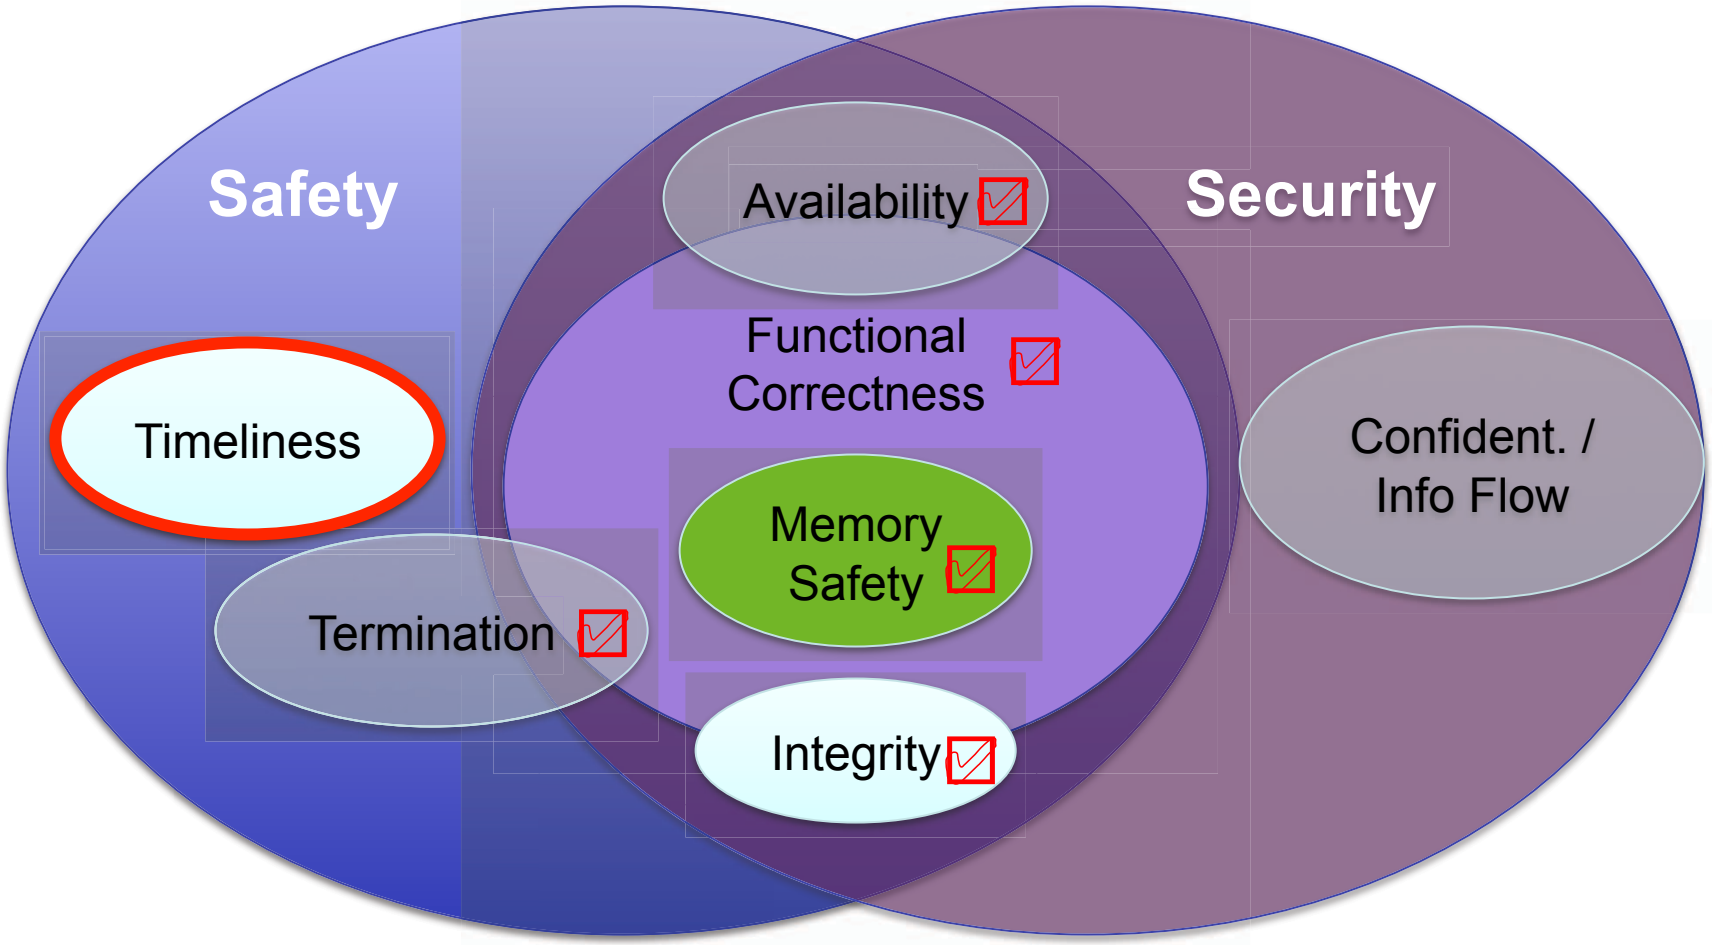

Timeliness

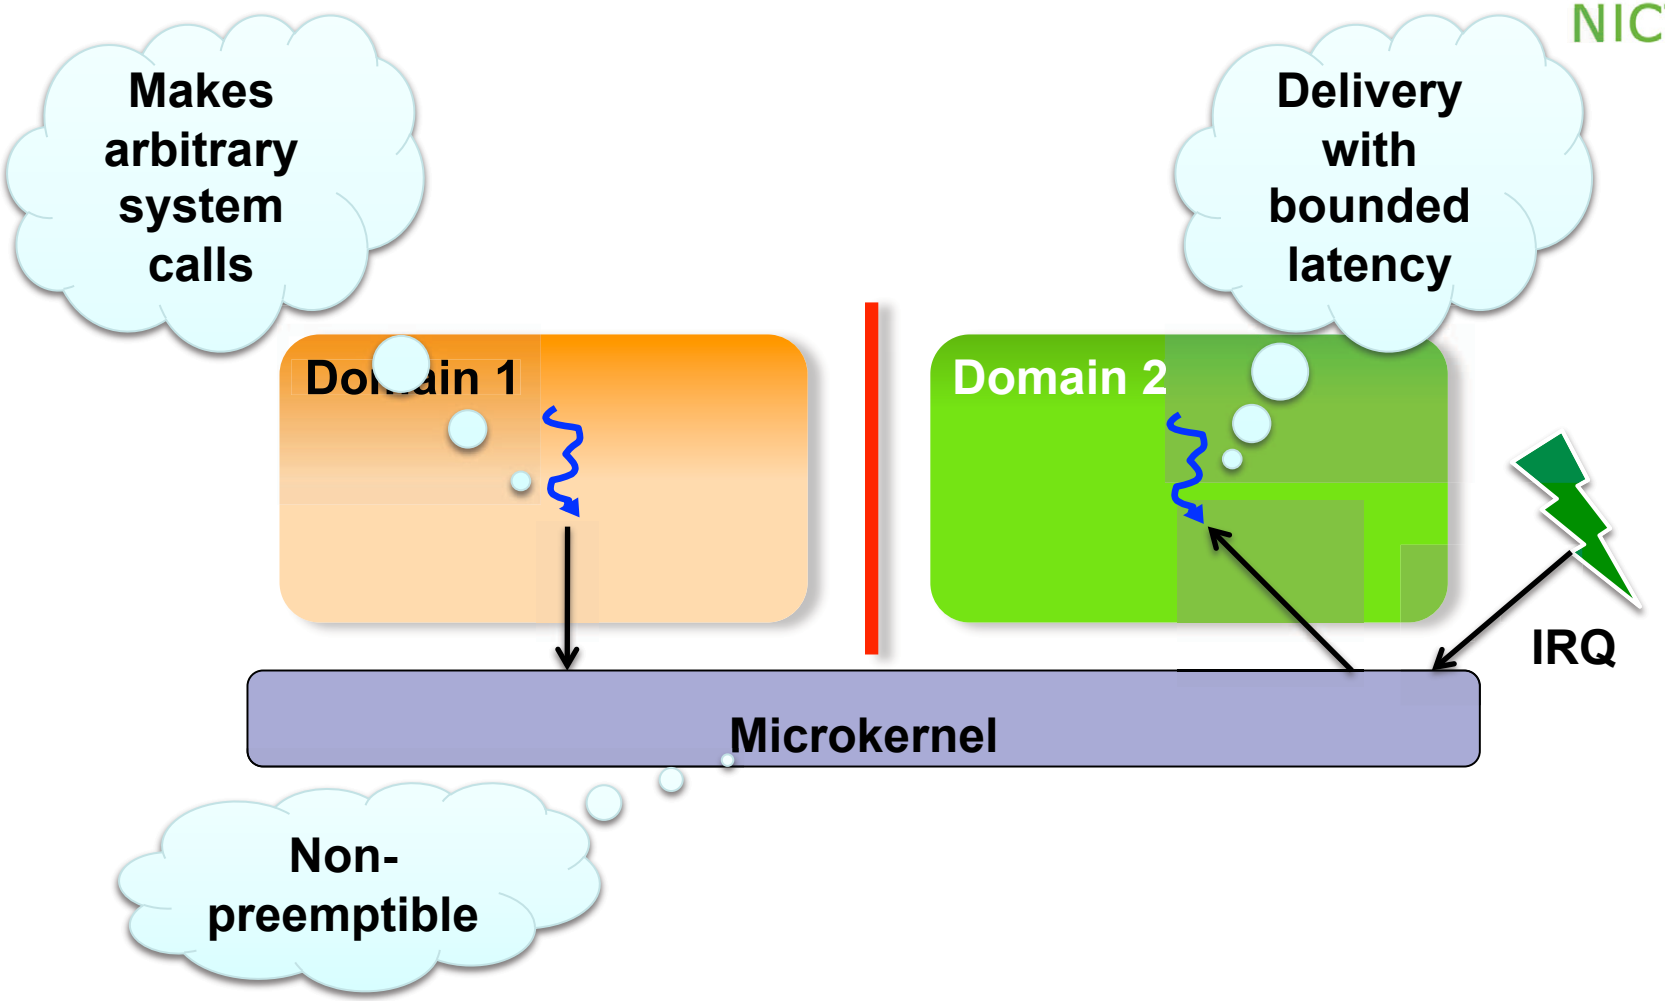

Need worst-case execution time (WCET) analysis of kernel

seL4 WCET Analysis Approach

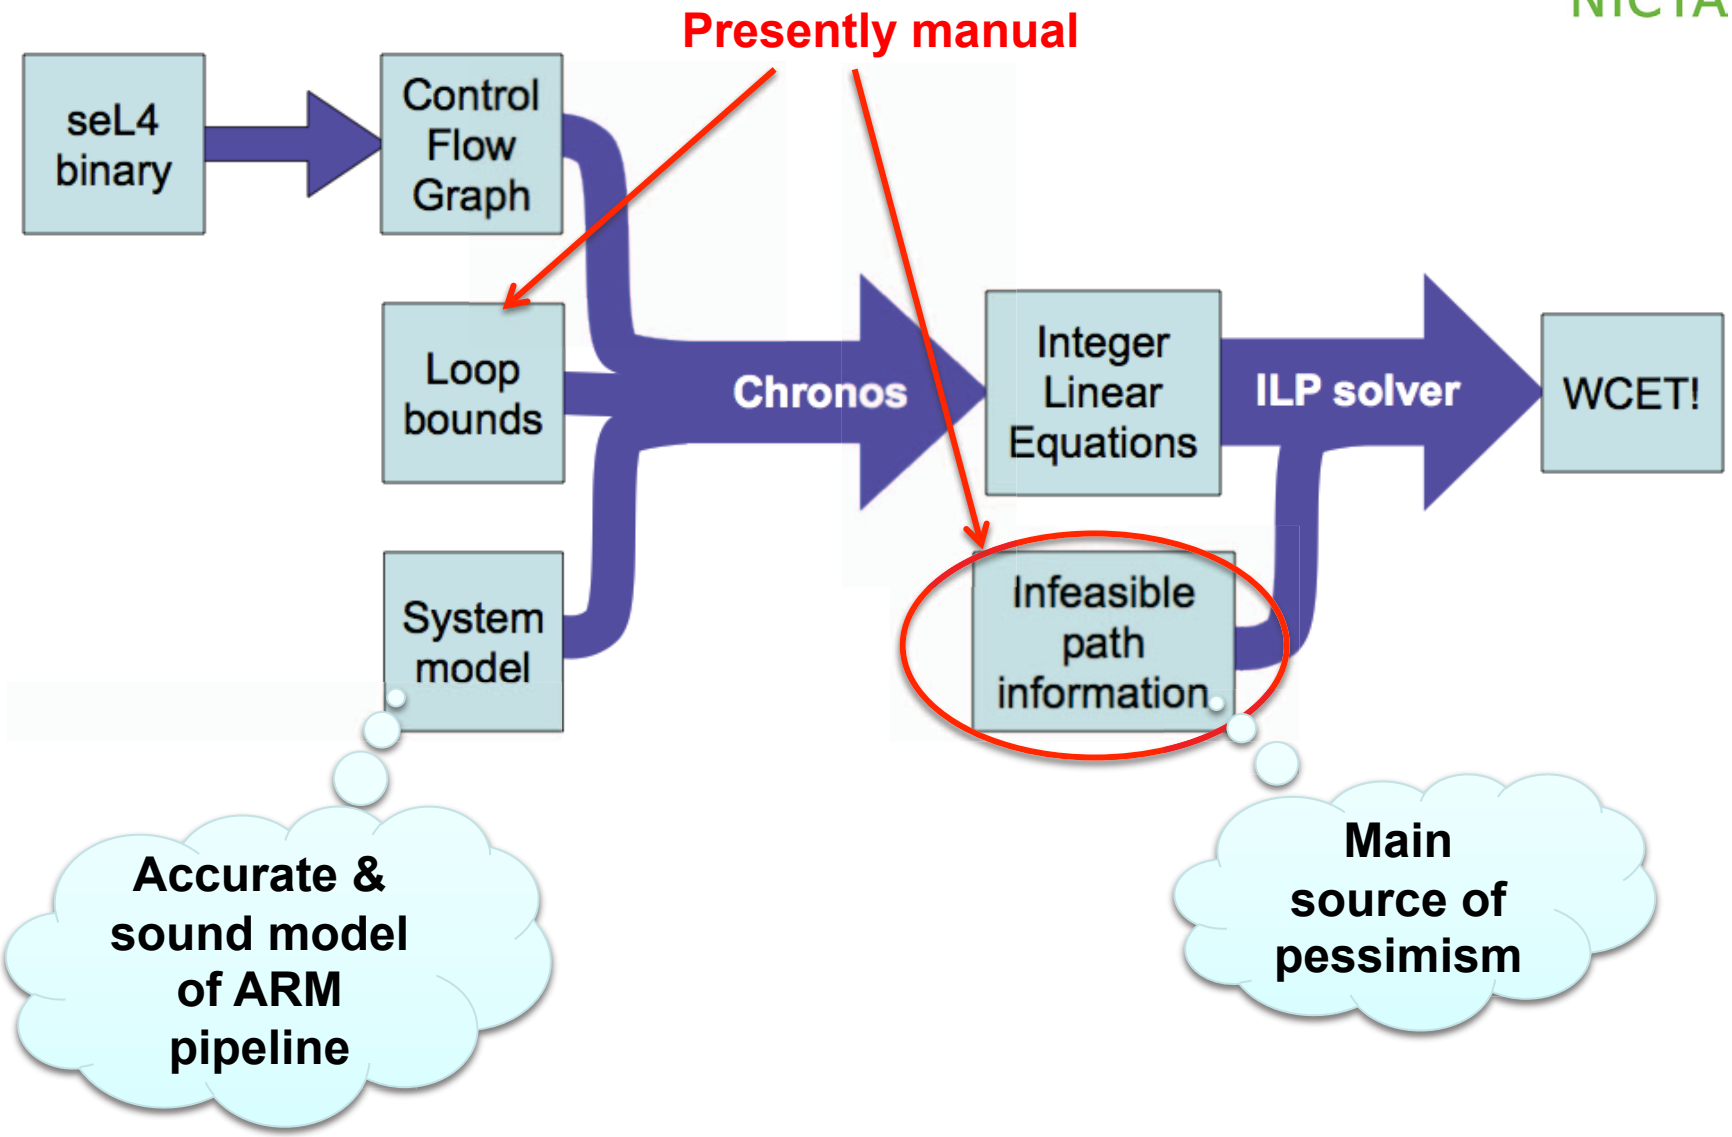

seL4 Worst-Case IRQ Latencies

First complete & sound WCET of a protected kernel [RTSS'11]

- Over 600 ms ☹

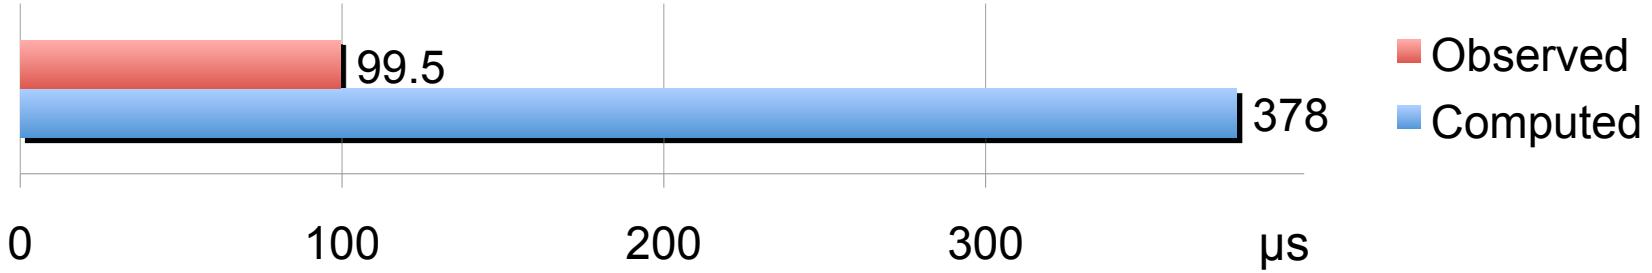

- Since improved by factor 1,500 [EuroSys'12]
  - Manual elimination of infeasible paths
  - Design and implementation changes, more is possible
  - Remaining pessimism is inevitable due to undefined HW behaviour

Future:

- Leverage verification invariants for loop bounds, infeasible paths
- Use as input for whole-system timing/schedulability analysis

seL4 for Safety and Security

NICTA

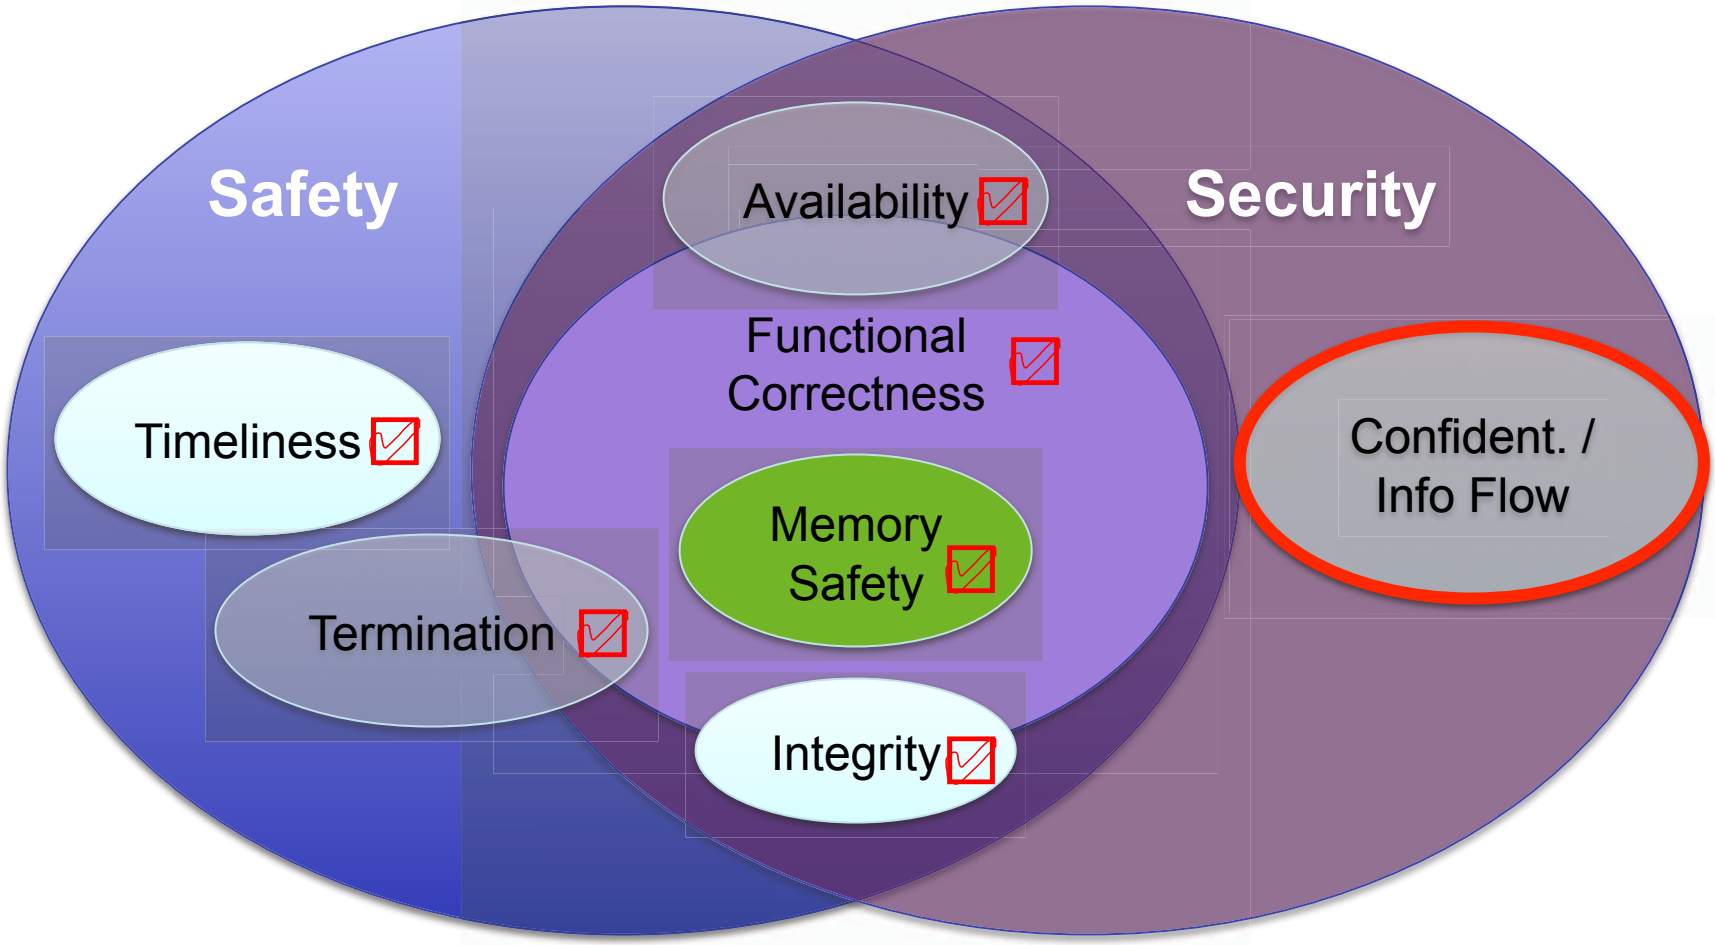

Confidentiality is about Read Accesses

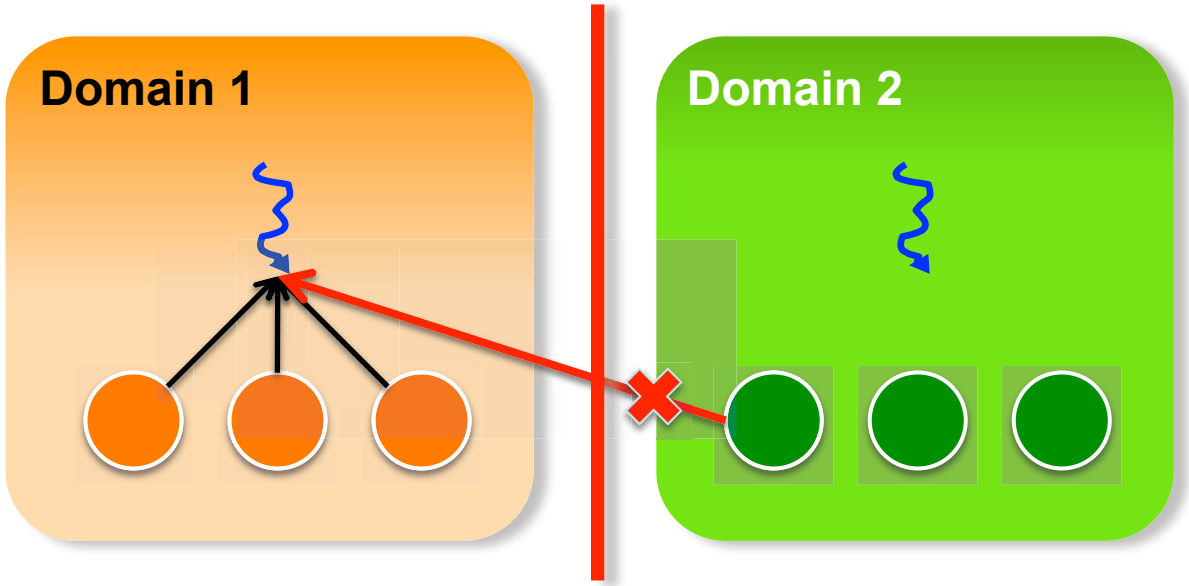

To prove:

- Domain-1 doesn't have read capabilities to Domain-2 objects
  ☐ no action of any agents will reveal Domain-2 state to Domain-1
- Harder than write, as protected data doesn't change
  - Violation not observable in Domain-2!
- Use non-interference: Domain-2 execution cannot affect Domain-1
- In progress!

Covert Channels?

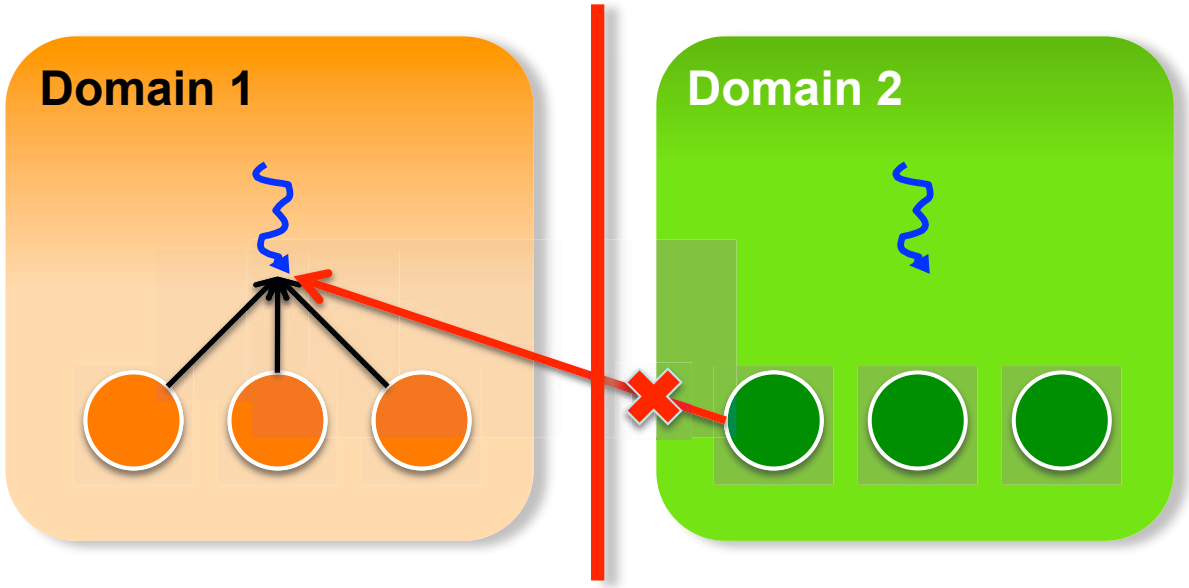

Storage channels:

- Should be able to eliminate by non-interference
  - … but need low-level machine model

Timing channels:

- May be able to leverage WCET analysis techniques?
- Not even started yet…

seL4 for Safety and Security

NICTA

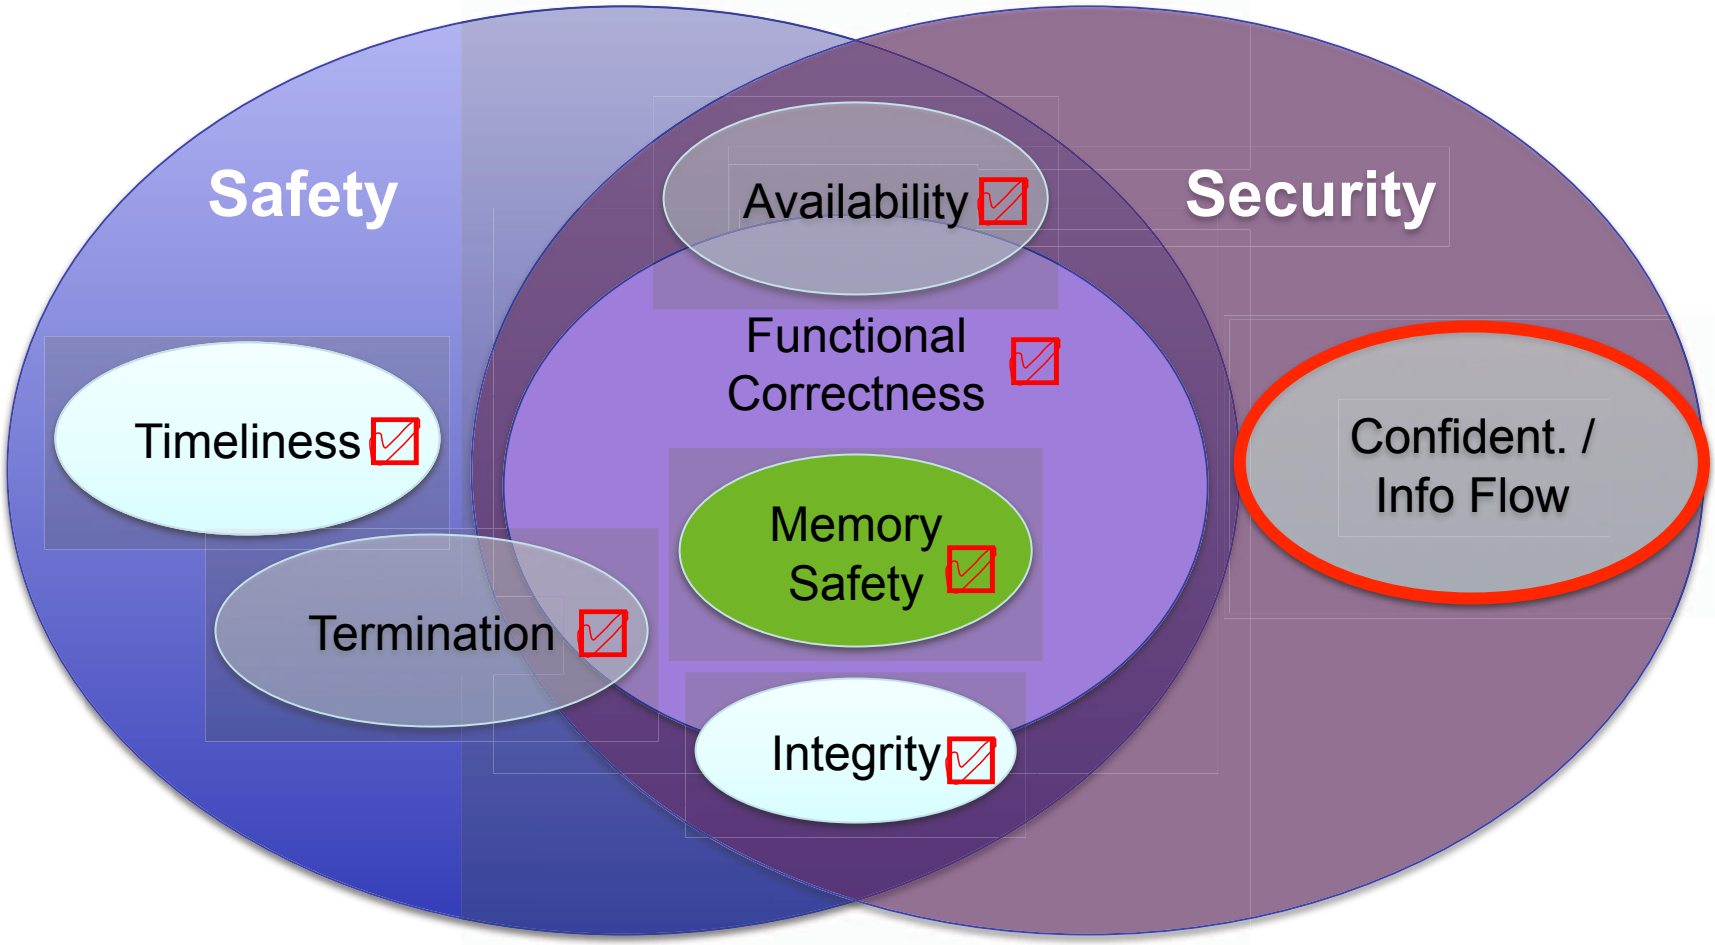

The seL4 Experience

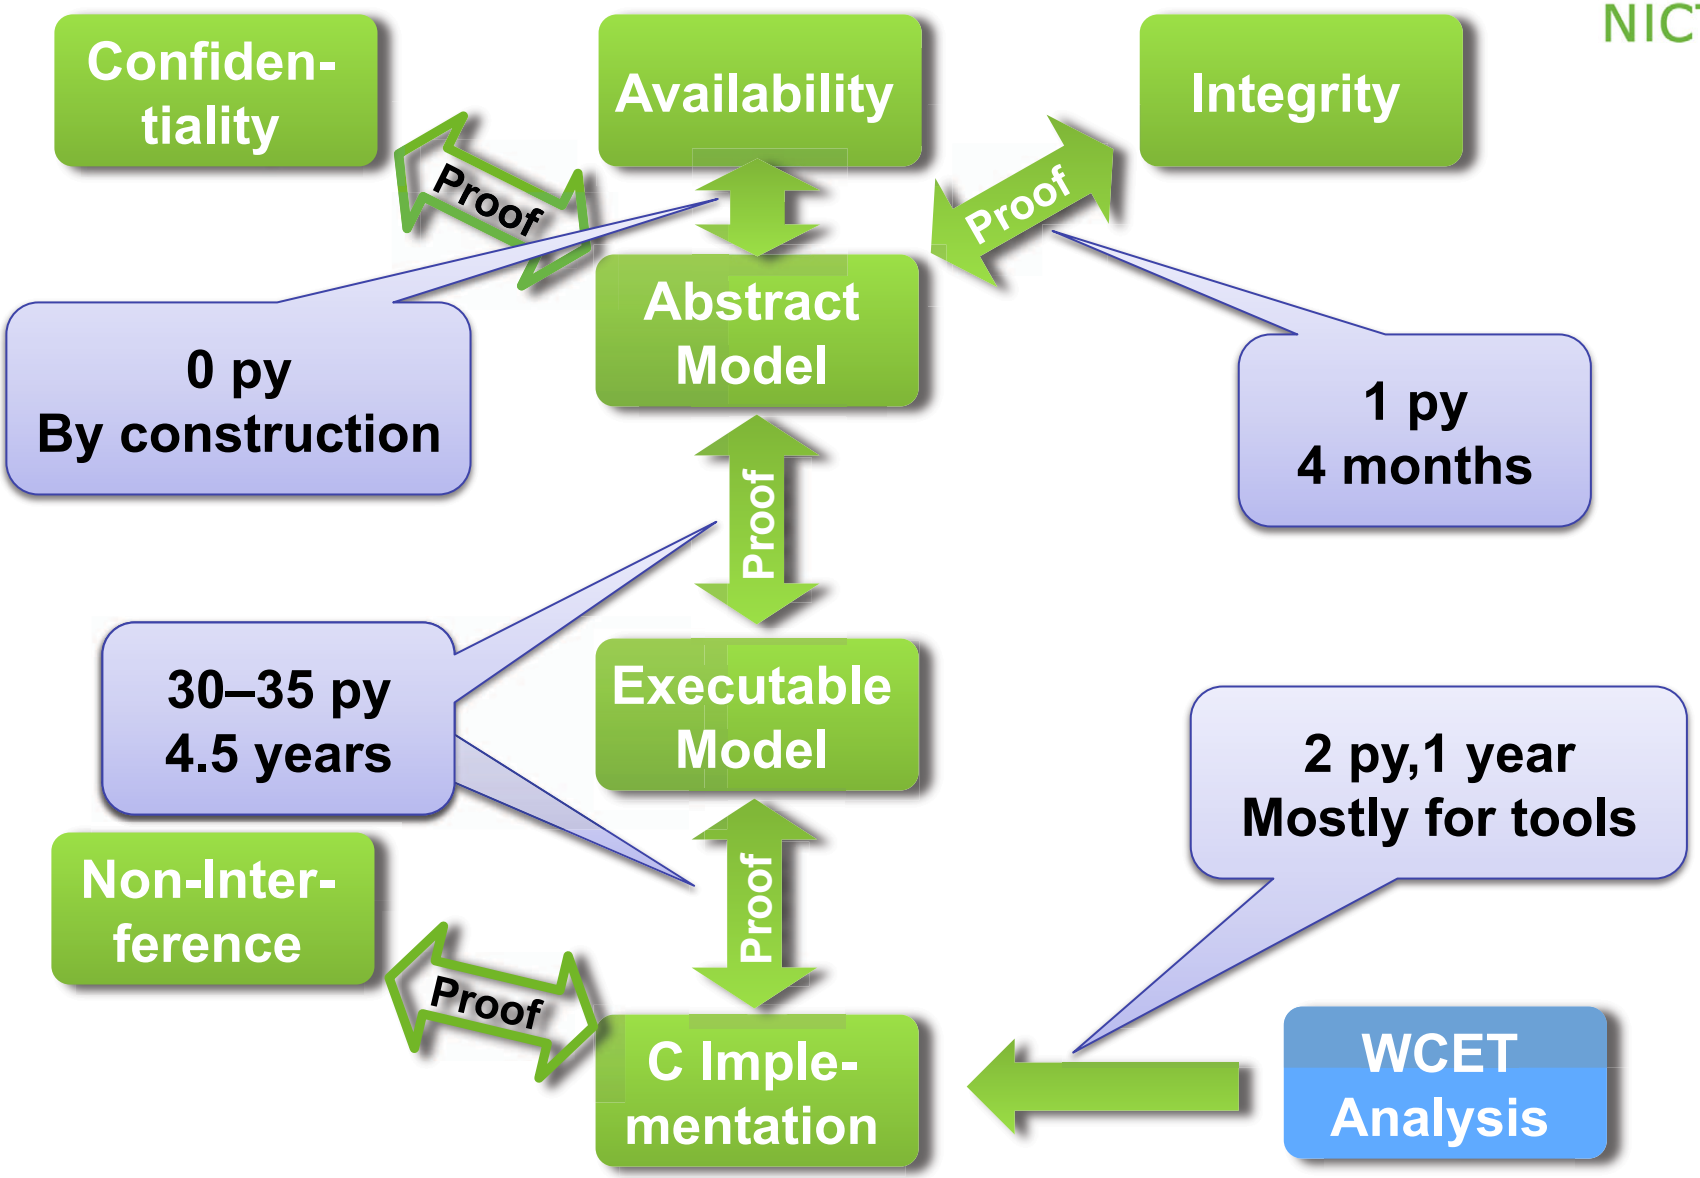

seL4: Next 12–24 Months

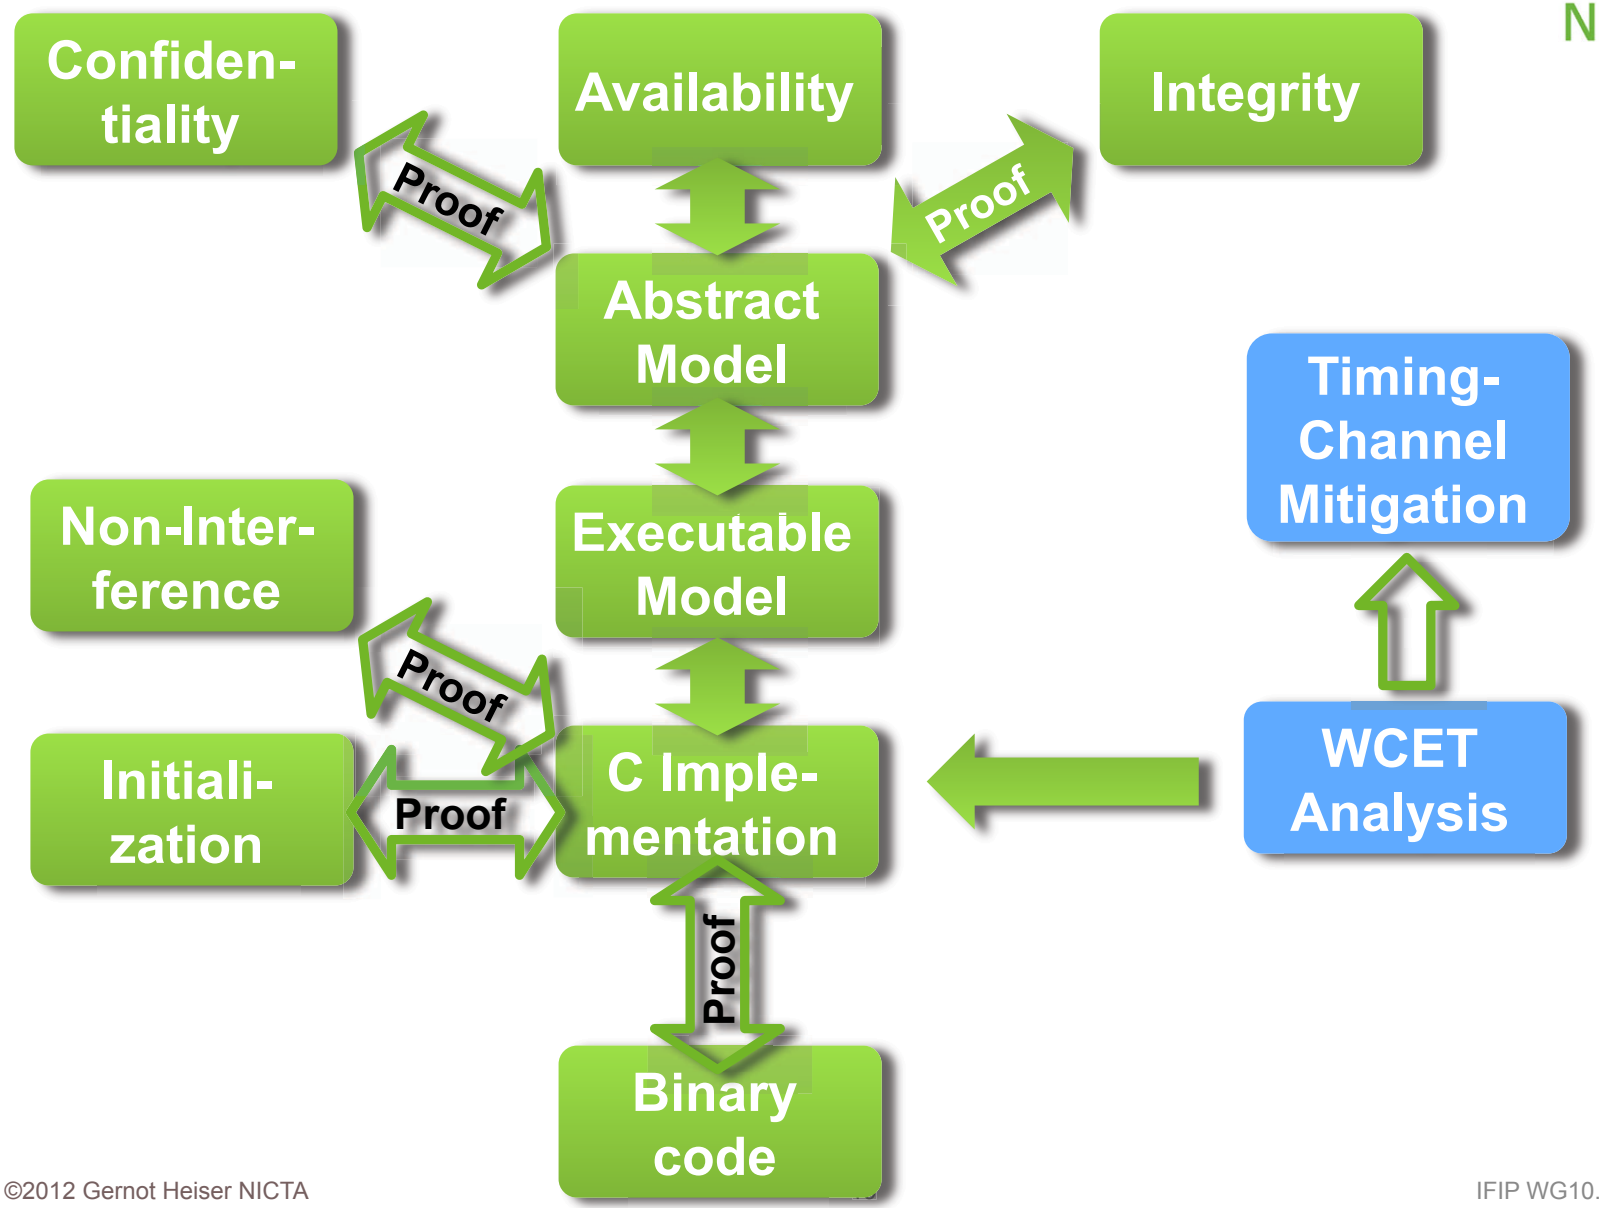

Binary Code Verification (In Progress)

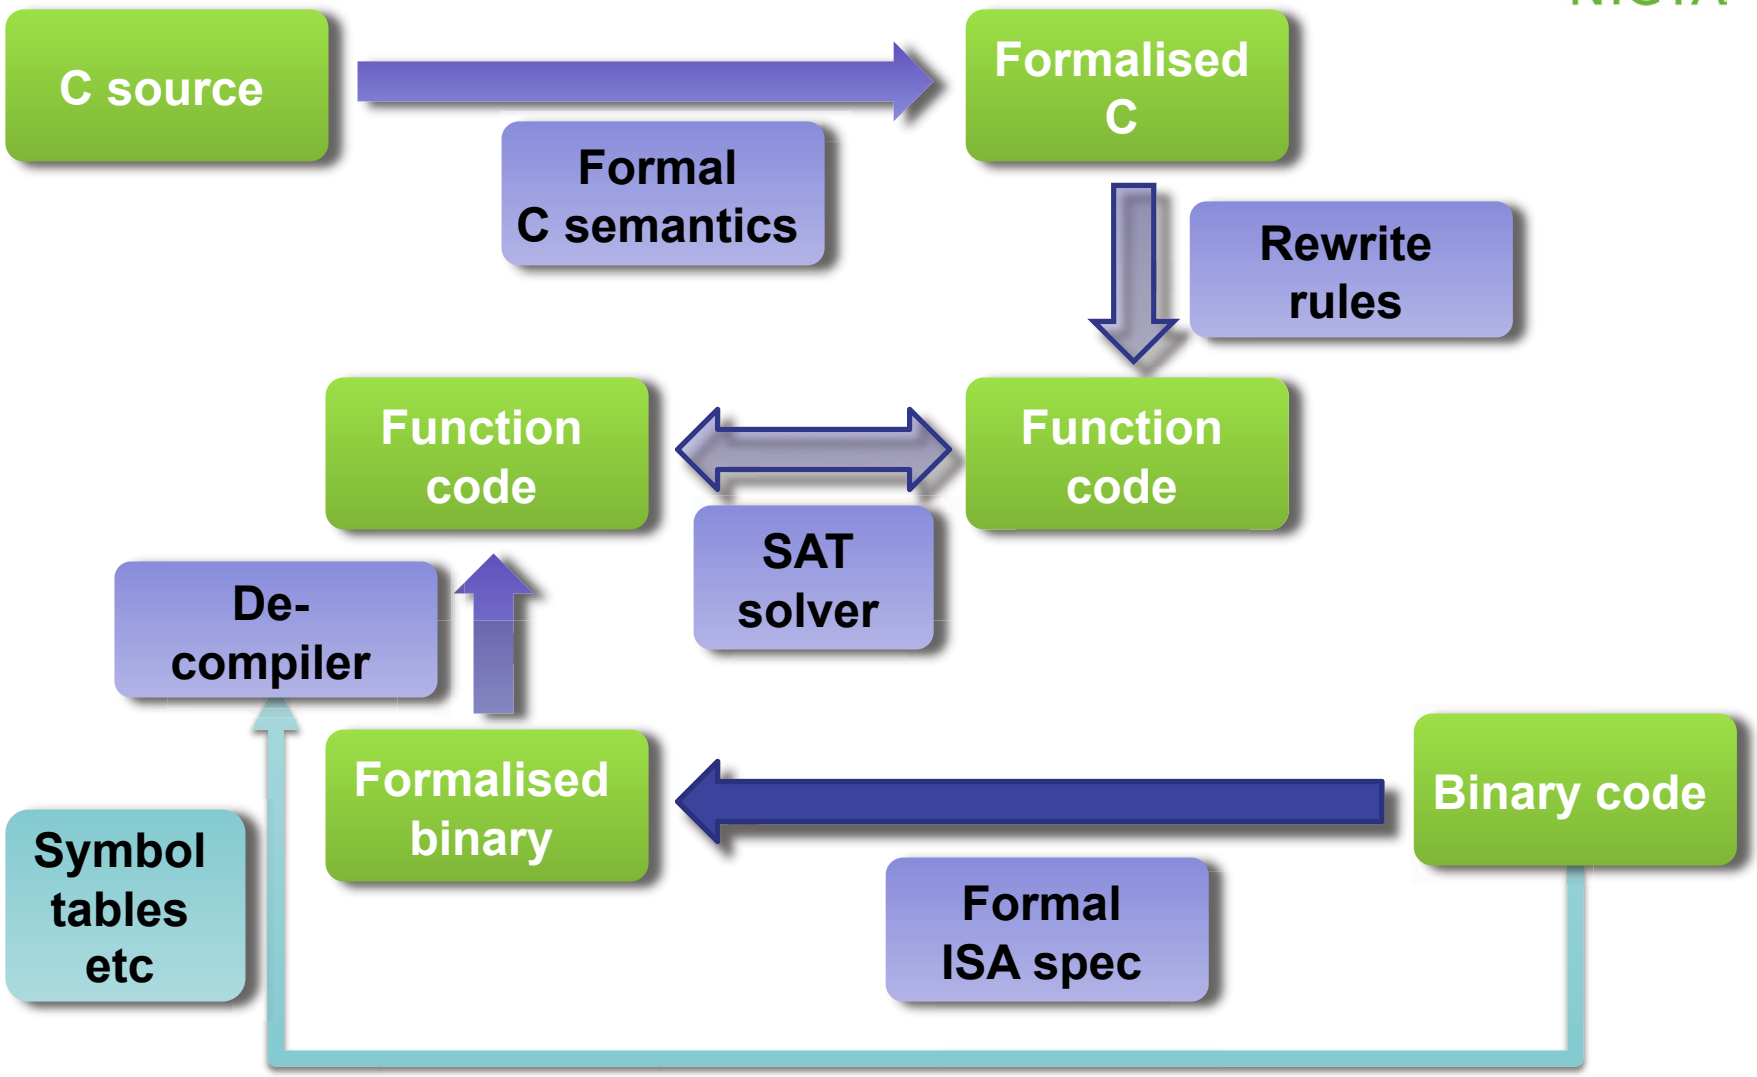

Long-Term View

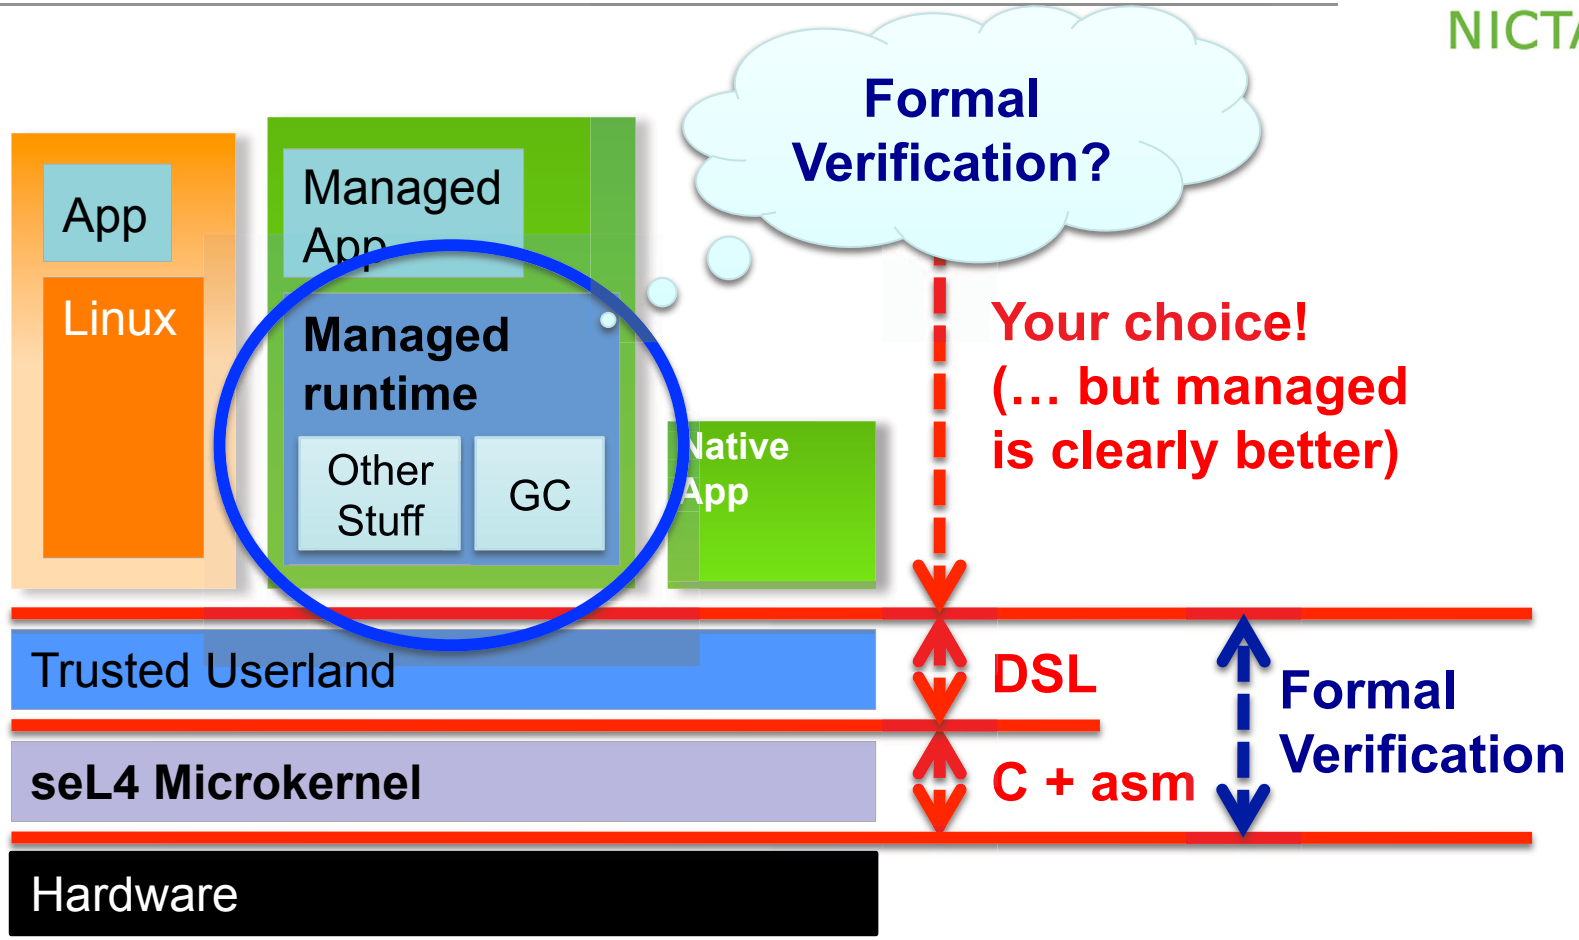

mailto:gernot@nicta.com.au

Google: “ertos”