

NO SAFETY WITHOUT SECURITY

Cybersecurity for Autonomous Vehicles

Gernot Heiser FTSA FIEEE FACM

Scientia Professor & John Lions Chair, UNSW Sydney

Chief Research Scientist, Data61, CSIRO

Cybersecurity: 1st Class Safety Issue



Fundamental rules of cyber space:

1. The internet is a hostile environment
2. Anything that is internet-connected **can** be attacked
3. Anything that **can** be attacked **will** be attacked

Examples:

- Cars
- Trains
- Aircraft

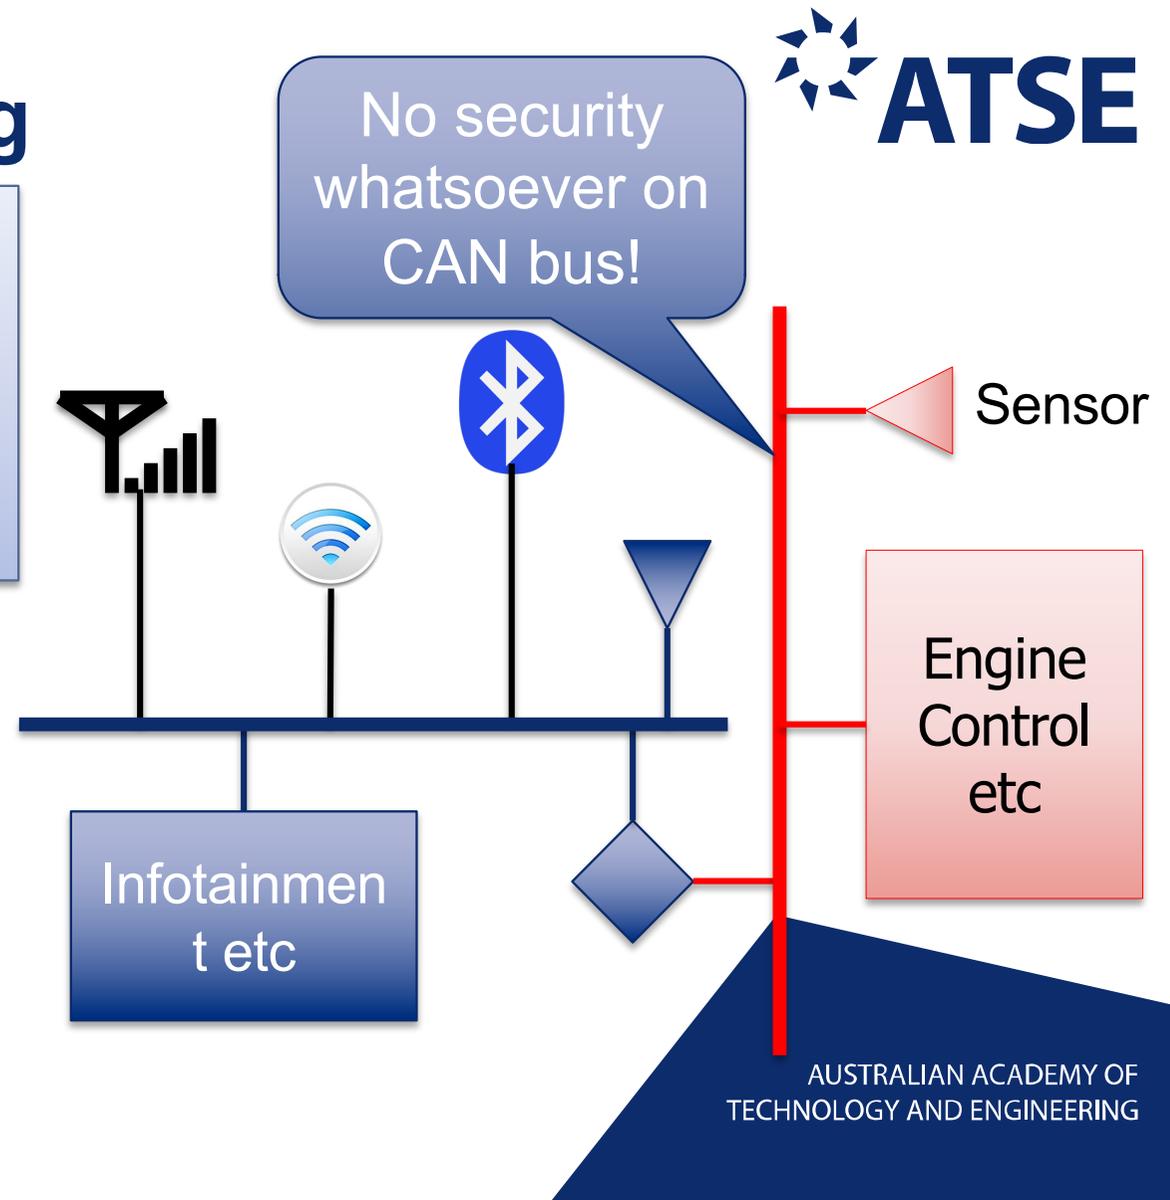
Example Car Hacking

Networking for:

- Entertainment
- Driver information
- Safety (tire pressure...)
- Maintenance (OTA upgrades)



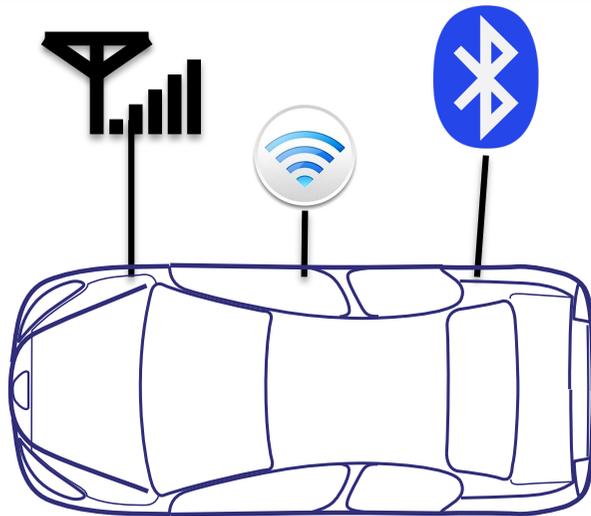
3 Heiser Korea 6'18



Challenge of Networking

Networking creates remote attack opportunities

- from passengers (wifi, Bluetooth)
- from nearby cars (wifi, Bluetooth) – incl infected ones!
- from anywhere (cellular)



Attack vectors:

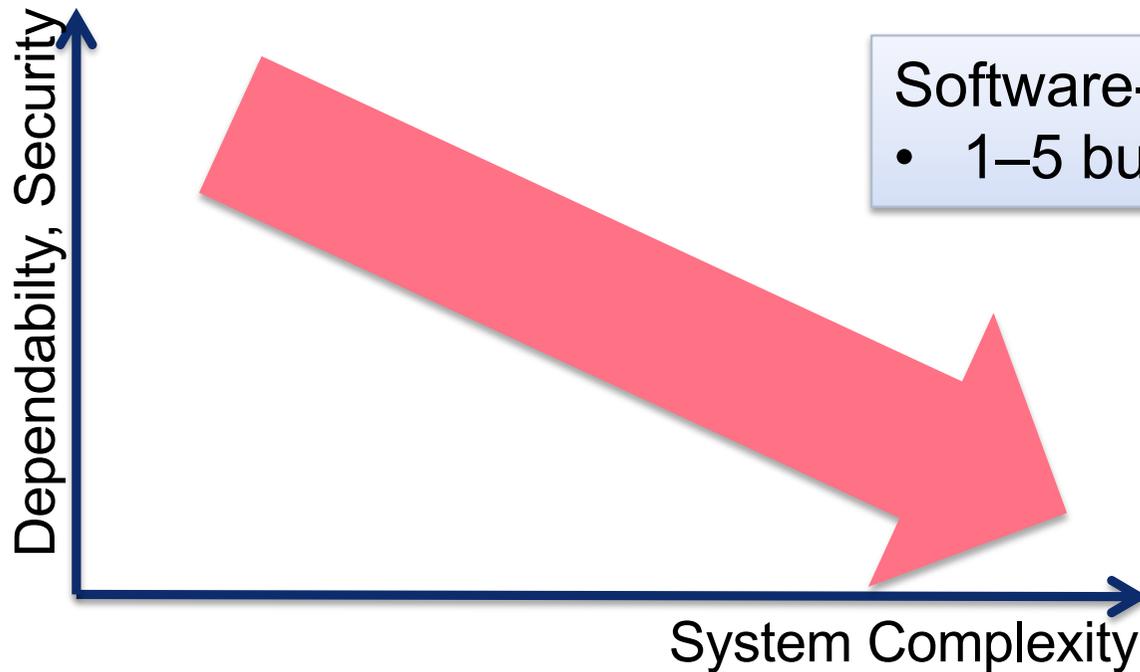
- Insecure protocols
- Reusing crypto keys
- **Software vulnerabilities**



BlueBorne

WHY ARE SYSTEMS SO VULNERABLE?

Failure Reason #1: Complexity



Software-engineering rule of thumb:

- 1–5 bugs per 1,000 lines of *quality* code

Bluetooth protocol stack:
Multiple 100,000 lines

Linux/Windows kernel:
Tens of millions lines

Complexity Drivers

- Features/functionality
- Legacy reuse

Operating-System “Security”

A screenshot of an Ars Technica article snippet. The top navigation bar is black with the "ars" logo in an orange circle, followed by "TECHNICA" and a search icon. Green navigation links include "BIZ & IT", "TECH", "SCIENCE", "POLICY", "CARS", and "GAMING". The article text is partially obscured by three red callout boxes. The visible text includes "RISK ASSESSMENT —", "Unsafe at any clock speed:", "Linux kernel security", "rethink", and "Ars reports from the Linux Security Summit—and finds much work that needs to be done." The red callout boxes contain the text: "Windows is no better!", "Software will break", and "The enemy will be on the platform!".

ars TECHNICA 🔍 BIZ & IT TECH SCIENCE POLICY CARS GAMING

RISK ASSESSMENT —

Windows is no better!

Unsafe at any clock speed:

Linux kernel security **Software will break**

rethink **The enemy will be on the platform!**

Ars reports from the Linux Security Summit—and finds much work that needs to be done.

Autonomy Increases Complexity

Requires new functionality

- Core autonomy functionality
- Computer vision
- Sensor fusion, from vastly increased number of sensors
- Collision avoidance

Increased integration of automotive control with external world

Much increased attack surface!

Failure Reason #2: Care Factor

Developer priorities

1. Features/functionality
2. Cost
3. Time to market
4. ...
5. ...
6. ...
7. ...
-
-
999. Security

Developer expertise

1. Undergraduate programming
2. Application domain
3. Maybe hardware
4. ...
5. ...
6. ...
7. ...
-
-
999. Security

Failure Reason #3: Security ≠ Safety



Classic safety thinking (eg automotive, avionics, electrical):

- Failures are *random*
- Failure rates can be kept *very low* through systematic process
- Multiple failures are *independent*

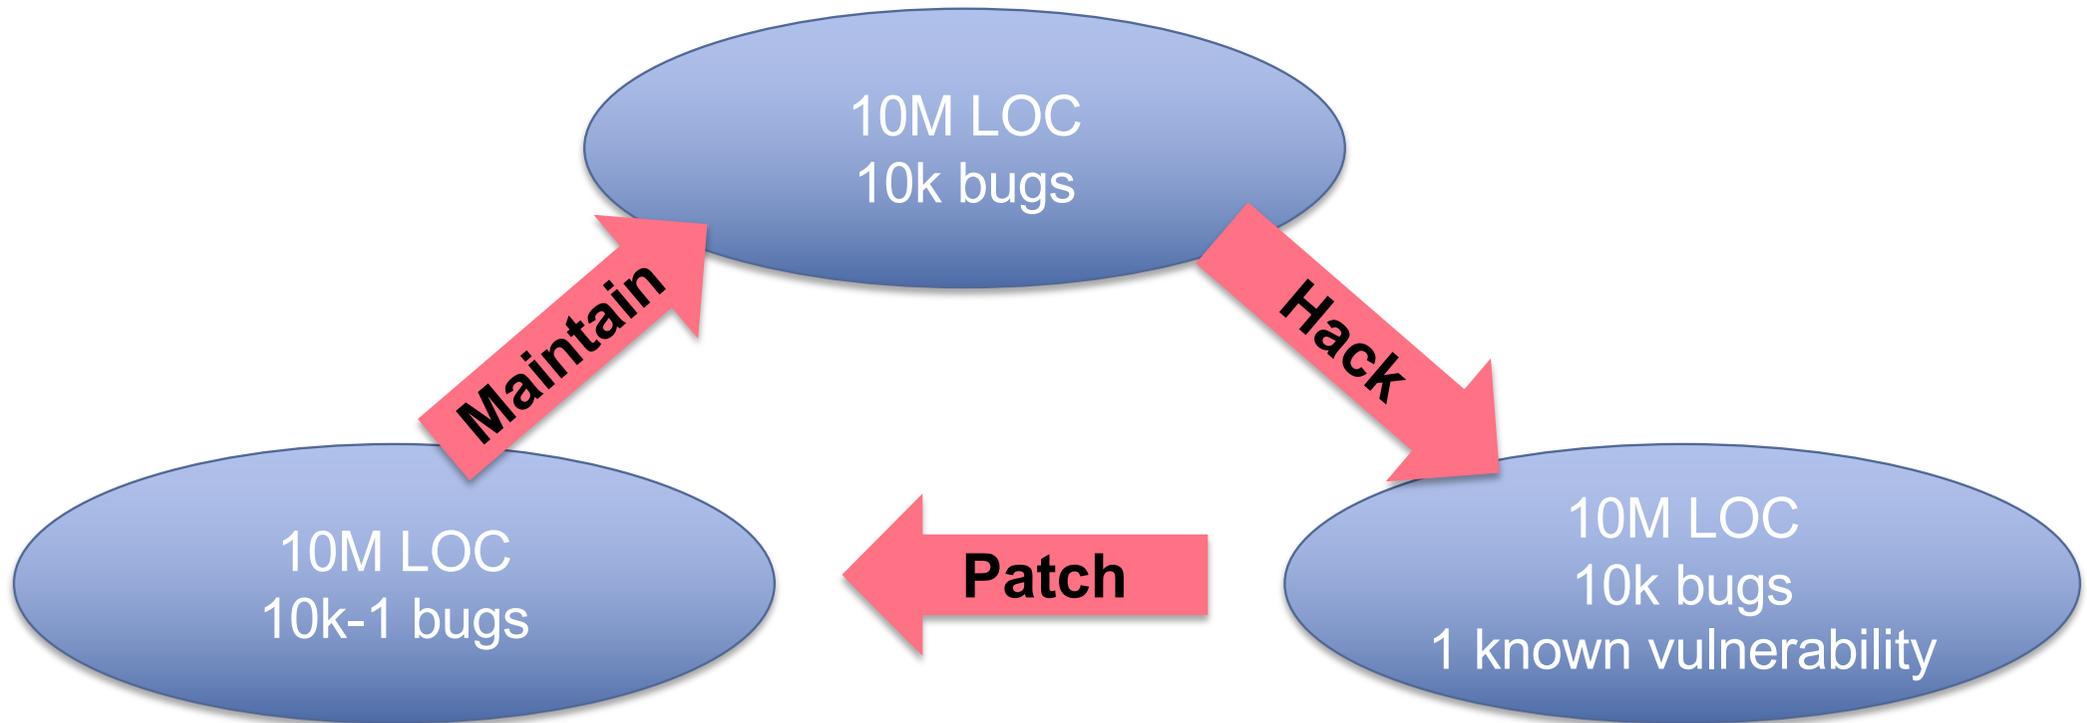
Reality of software security weaknesses:

- Failure is *deterministic*
- Failure rates are *high*
- Attackers *systematically combine* multiple vulnerabilities
- **Classic safety approaches do not work against cyber attacks!**

**No safety
without security!**

STANDARD DEFENCES

OK, So Let's Patch Regularly

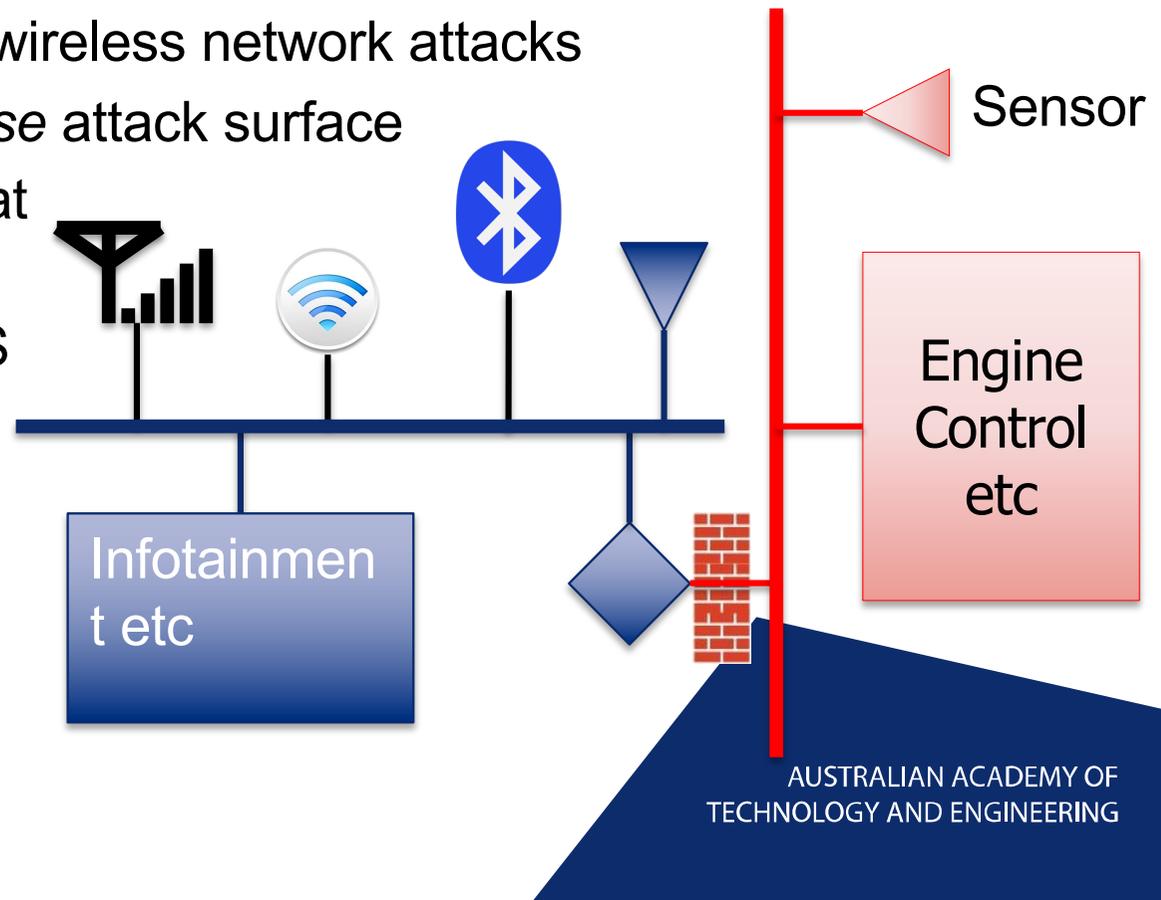


Patch-and-Pray: A losing proposition

So, Let's Use Firewalls!

- Imposes overhead (SWaP)
- Doesn't protect against edge, wireless network attacks
- Even more code – may *increase* attack surface
- No help for valid messages that trigger bugs in software
- Firewall runs on vulnerable OS

Firewalls treat symptoms,
not causes of problems,
are just another arms race!

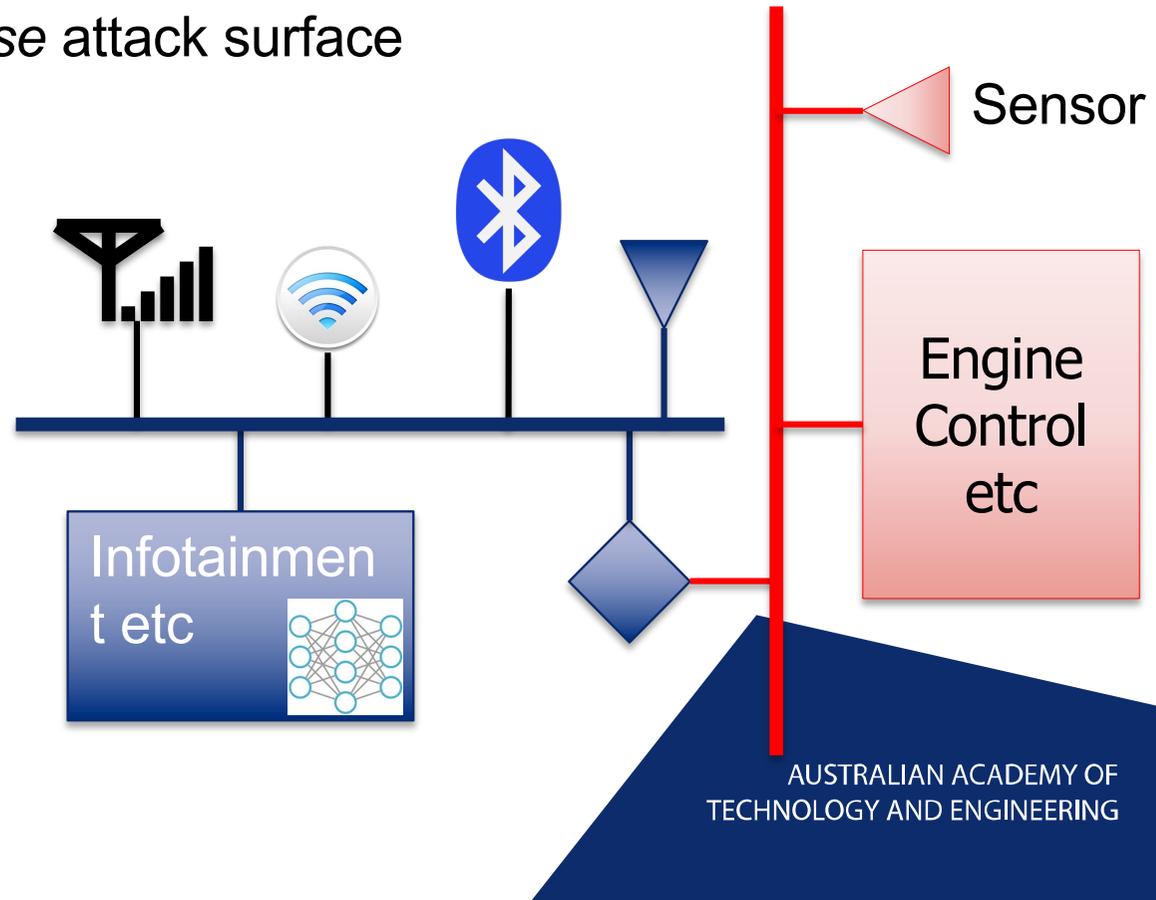


Let's Use AI to Detect Compromise!



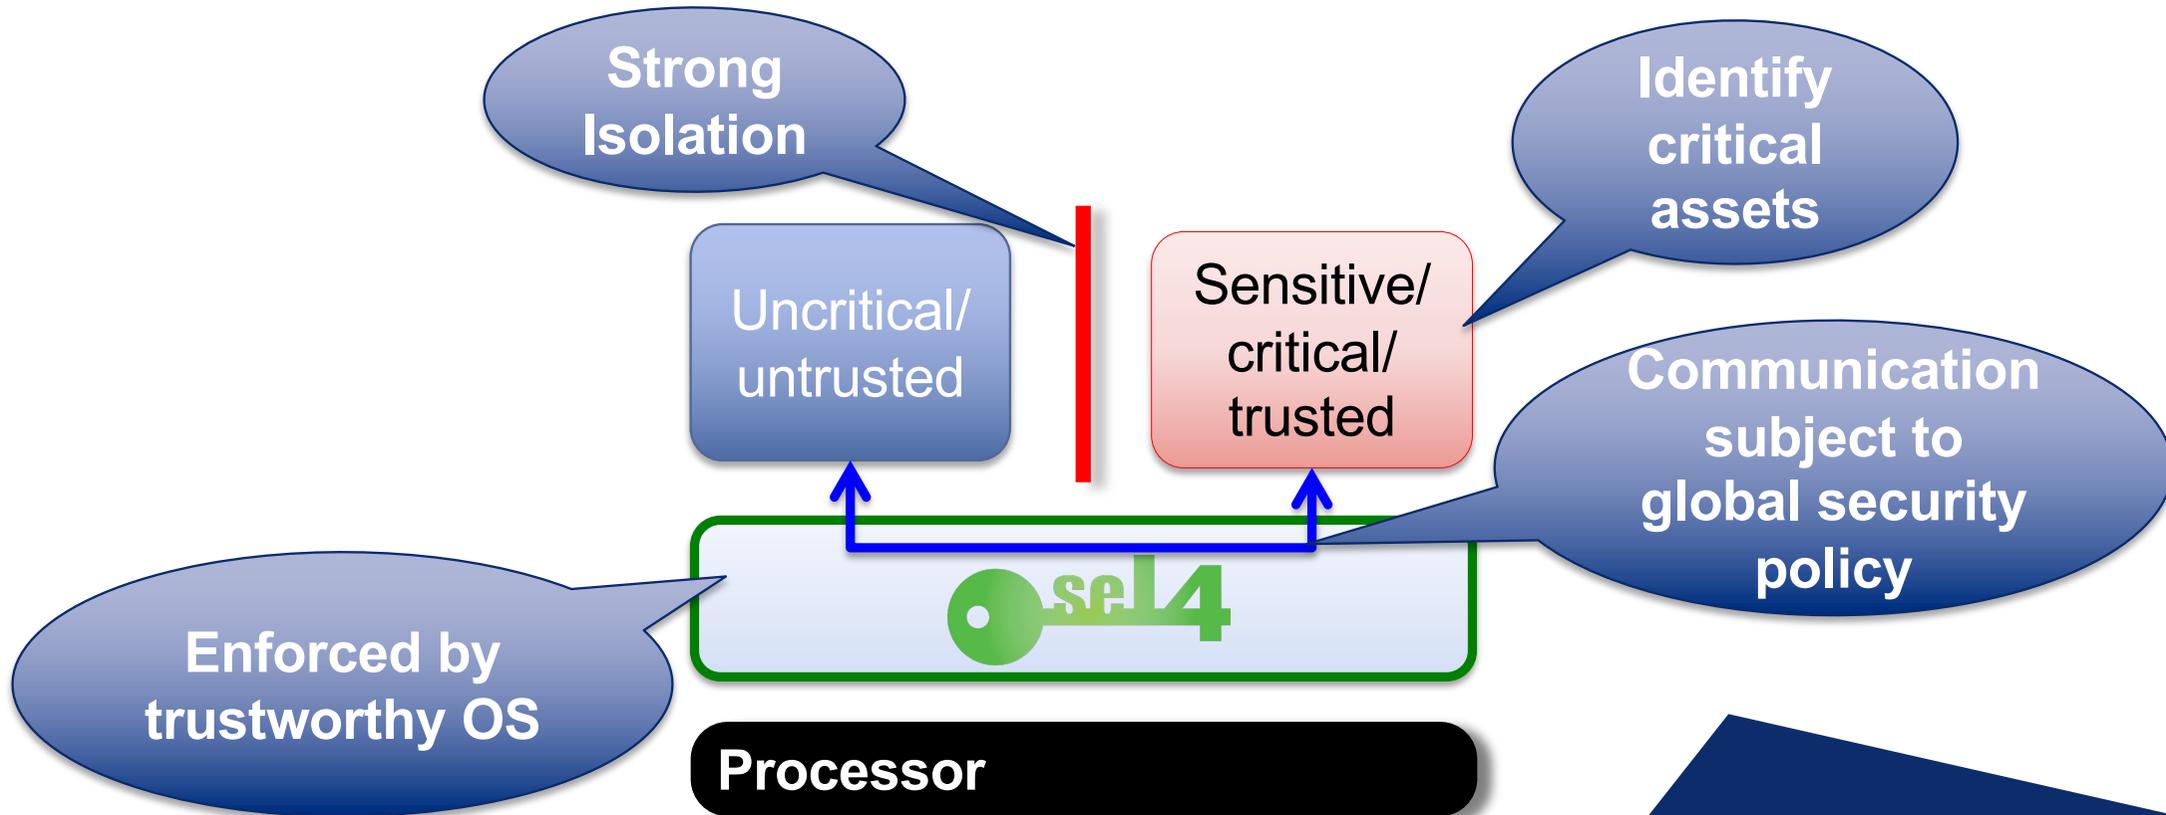
- Can only detect that system is already compromised
- Even more code – may *increase* attack surface
- Runs on compromised OS!

Intrusion detection –
admission of defeat



WHAT IS NEEDED?

Fundamental Requirement: Isolation



Trustworthiness: Can We Rely on Isolation?

A system is **trustworthy** if and only if:

- it behaves **exactly** as it is specified,
- in a **timely** manner, and
- while ensuring **secure** execution

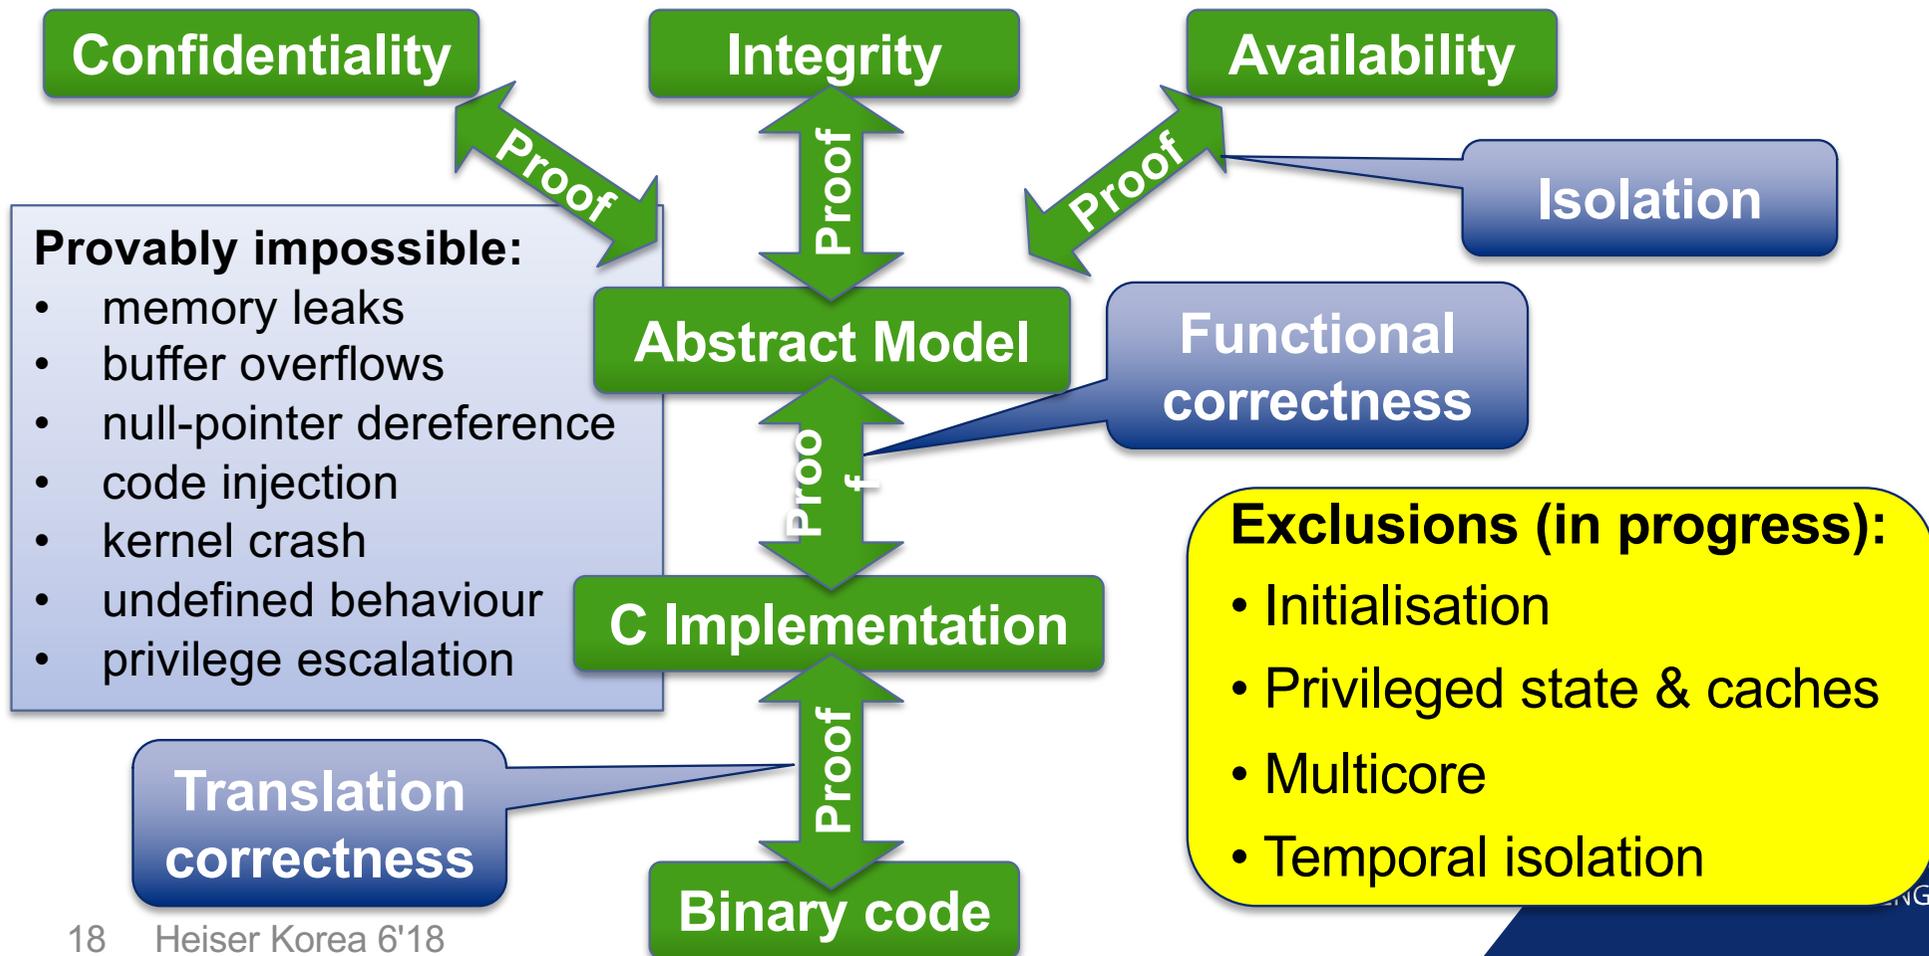
Claim:

A system must be considered *untrustworthy*, unless *proved* otherwise!

Corollary [with apologies to Dijkstra]:

Testing, code inspection, etc. can only show *lack of trustworthiness*!

seL4 Microkernel: We Have Proof!



How Does seL4 Compare?

Feature	seL4	Others
Performance	Fastest	5–10 × slower
Impl. bugs	Provably none	No guarantee
Isolation	Proved	No guarantee
Latency bounds	Sound and complete	Estimates only
Storage channels	Provably none	No guarantee
Timing channels	Low-overhead prevention	No story or high overhead
Mixed criticality	Supported with high utilisation	None or resource-wastive

“World’s most verified kernel”

“Software you can depend on, data access you can trust”

Security by Architecture



Cyber-retrofit!

Uncritical/
untrusted

Apps

Linux

Virtual
machine
for legacy

Incremental
process: migrate
in pieces

Extract
critical bits,
run native

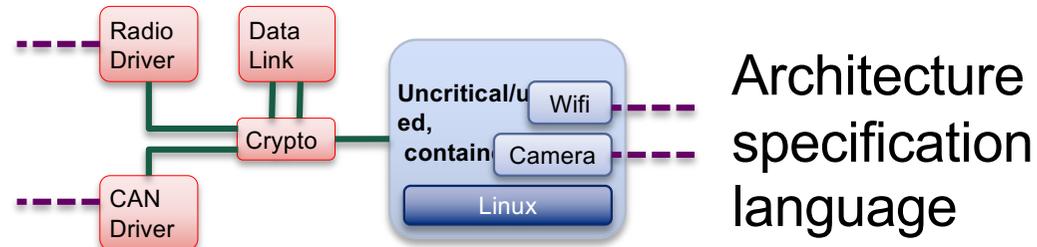
Critical
control

Device
driver

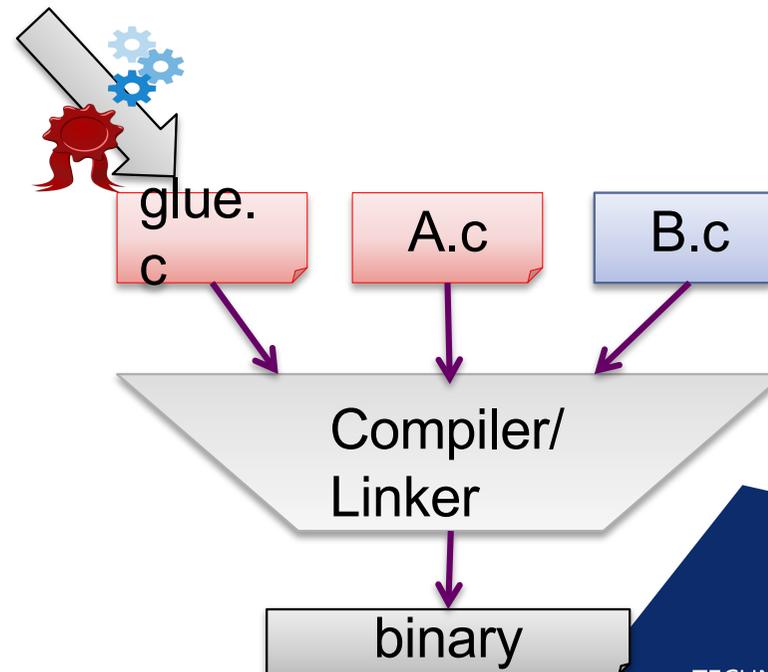
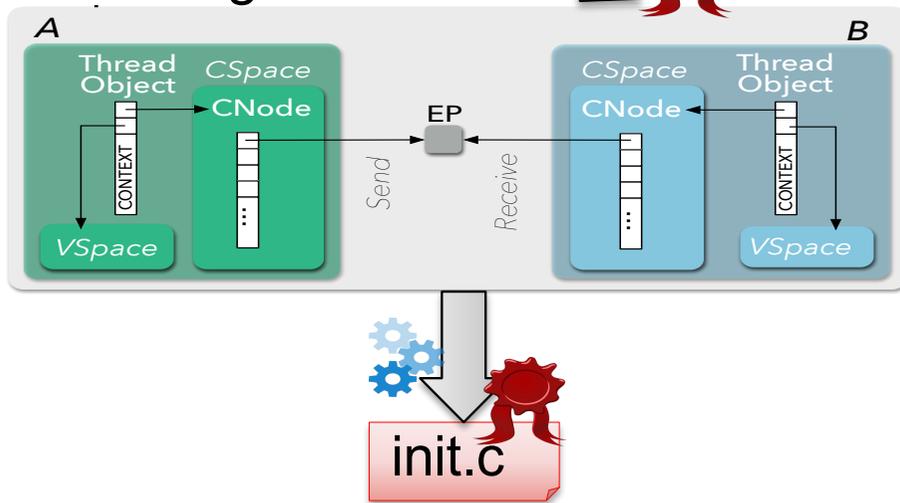
NW
stack



Enforcing the Architecture



Low-level access rights



Military-Grade Security



Cross-Domain
Desktop
Compositor



Boeing
Unmanned
Helicopter

US Army
Autonomous
Trucks



Crypto
Stick

Trustworthy Software At Work





<https://trustworthy.systems>

AUSTRALIAN ACADEMY OF
TECHNOLOGY AND ENGINEERING