

# Concurrent, Compositional Declassification without Reinventing the Wheel



THE UNIVERSITY OF  
MELBOURNE

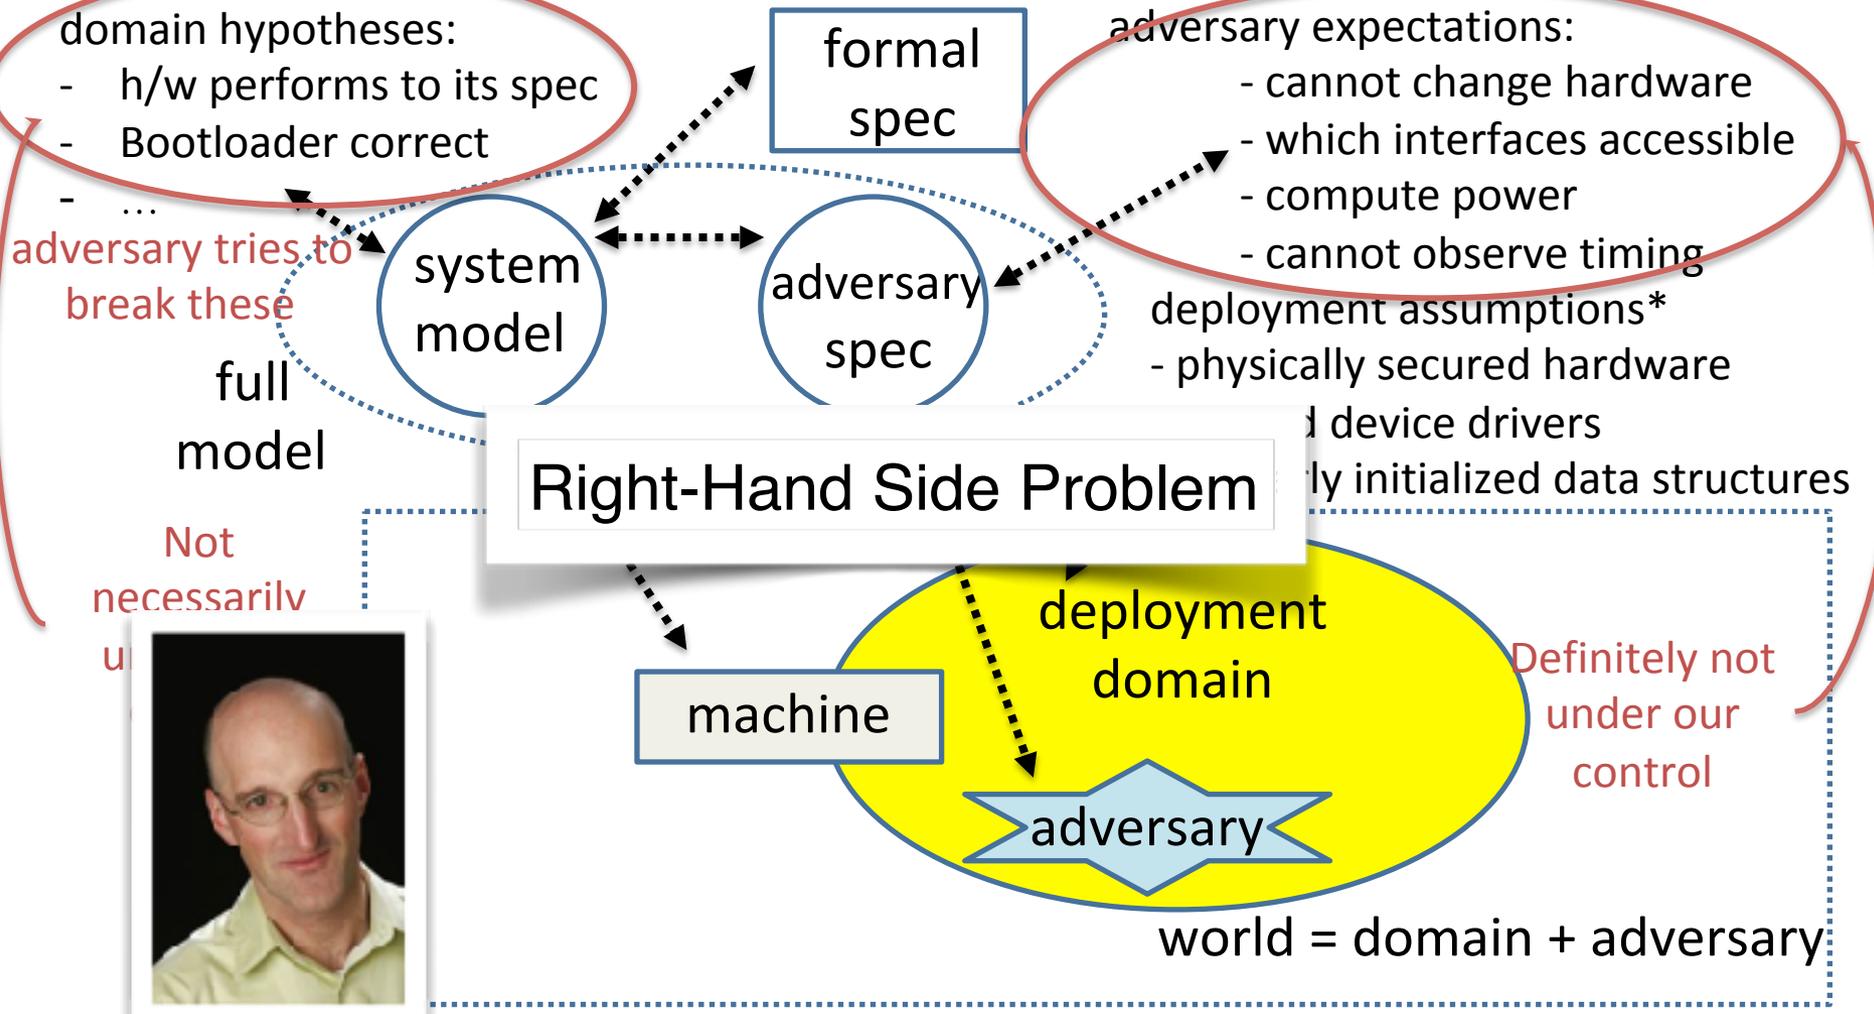
**Toby Murray**

University of Melbourne,  
Data61 (formerly NICTA), CSIRO

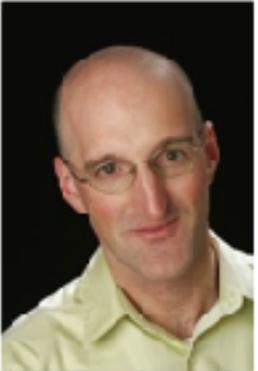
IFIP WG2.3, May 2018



# Models, Assumptions & World-Machine (OS)



Right-Hand Side Problem



Paul van Oorschot

# Confidentiality (Information Flow) Proofs



CoCon

CoSMed



**SEQUENTIAL**



CERTIKOS

SAFE



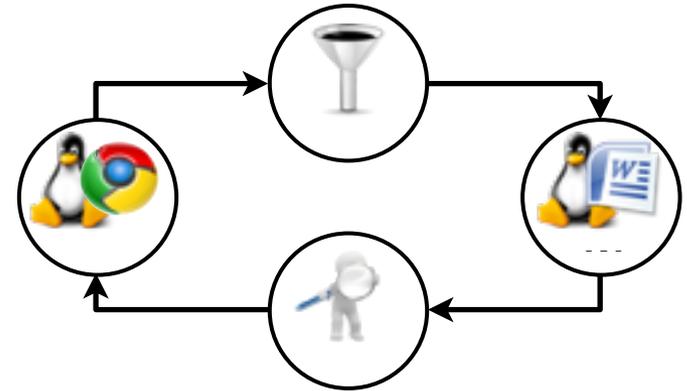
Quark



# Sequential vs Whole-System

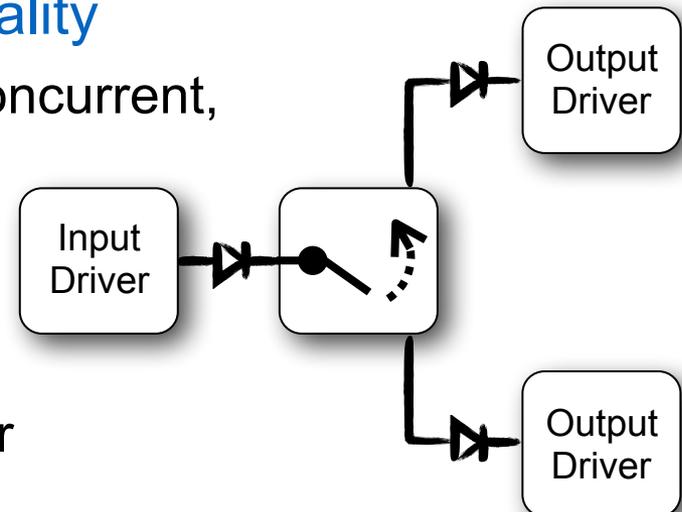
## Sequential Confidentiality

assumes application components are maximally hostile  
proved security only by reasoning about mandatory, enforced abstractions  
simple static policies only



## Concurrent, Whole-System Confidentiality

reasons about interactions between concurrent, application components  
proves that their composition yields a secure system  
runtime state-dependent and data-dependent policies and behaviour





Mark Beaumont



Kevin Elphinstone

# CDDC

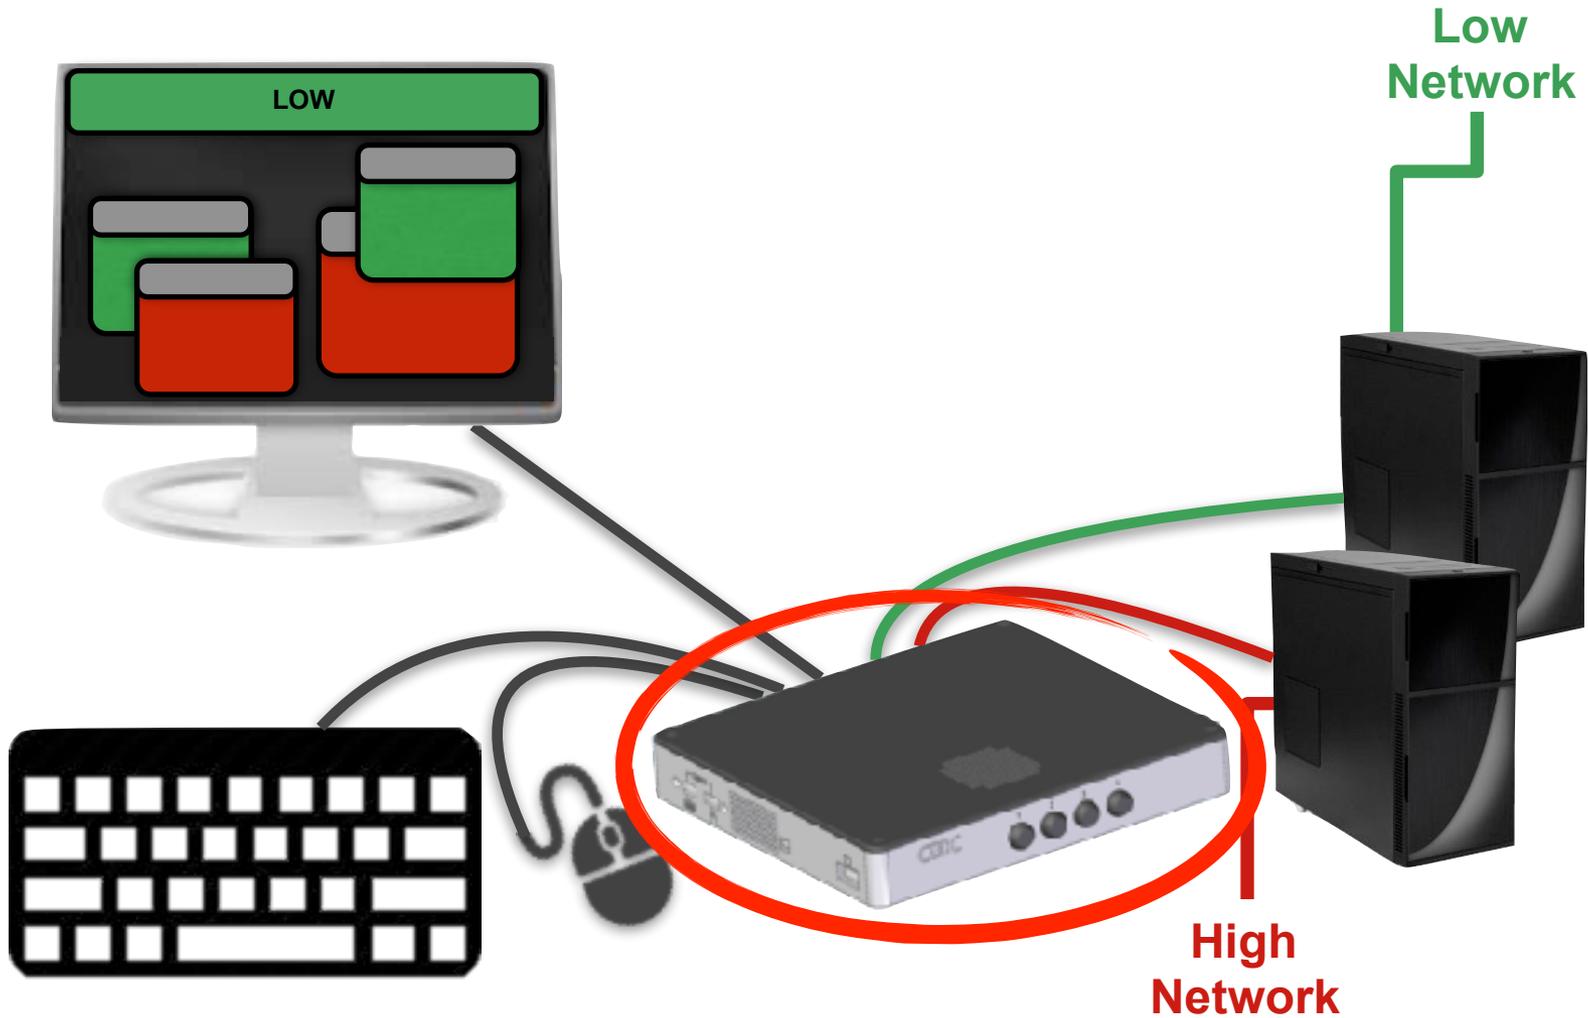
The Cross Domain Desktop Compositor



Australian Government  
Department of Defence  
Defence Science and Technology Group



# Cross Domain Desktop Compositor

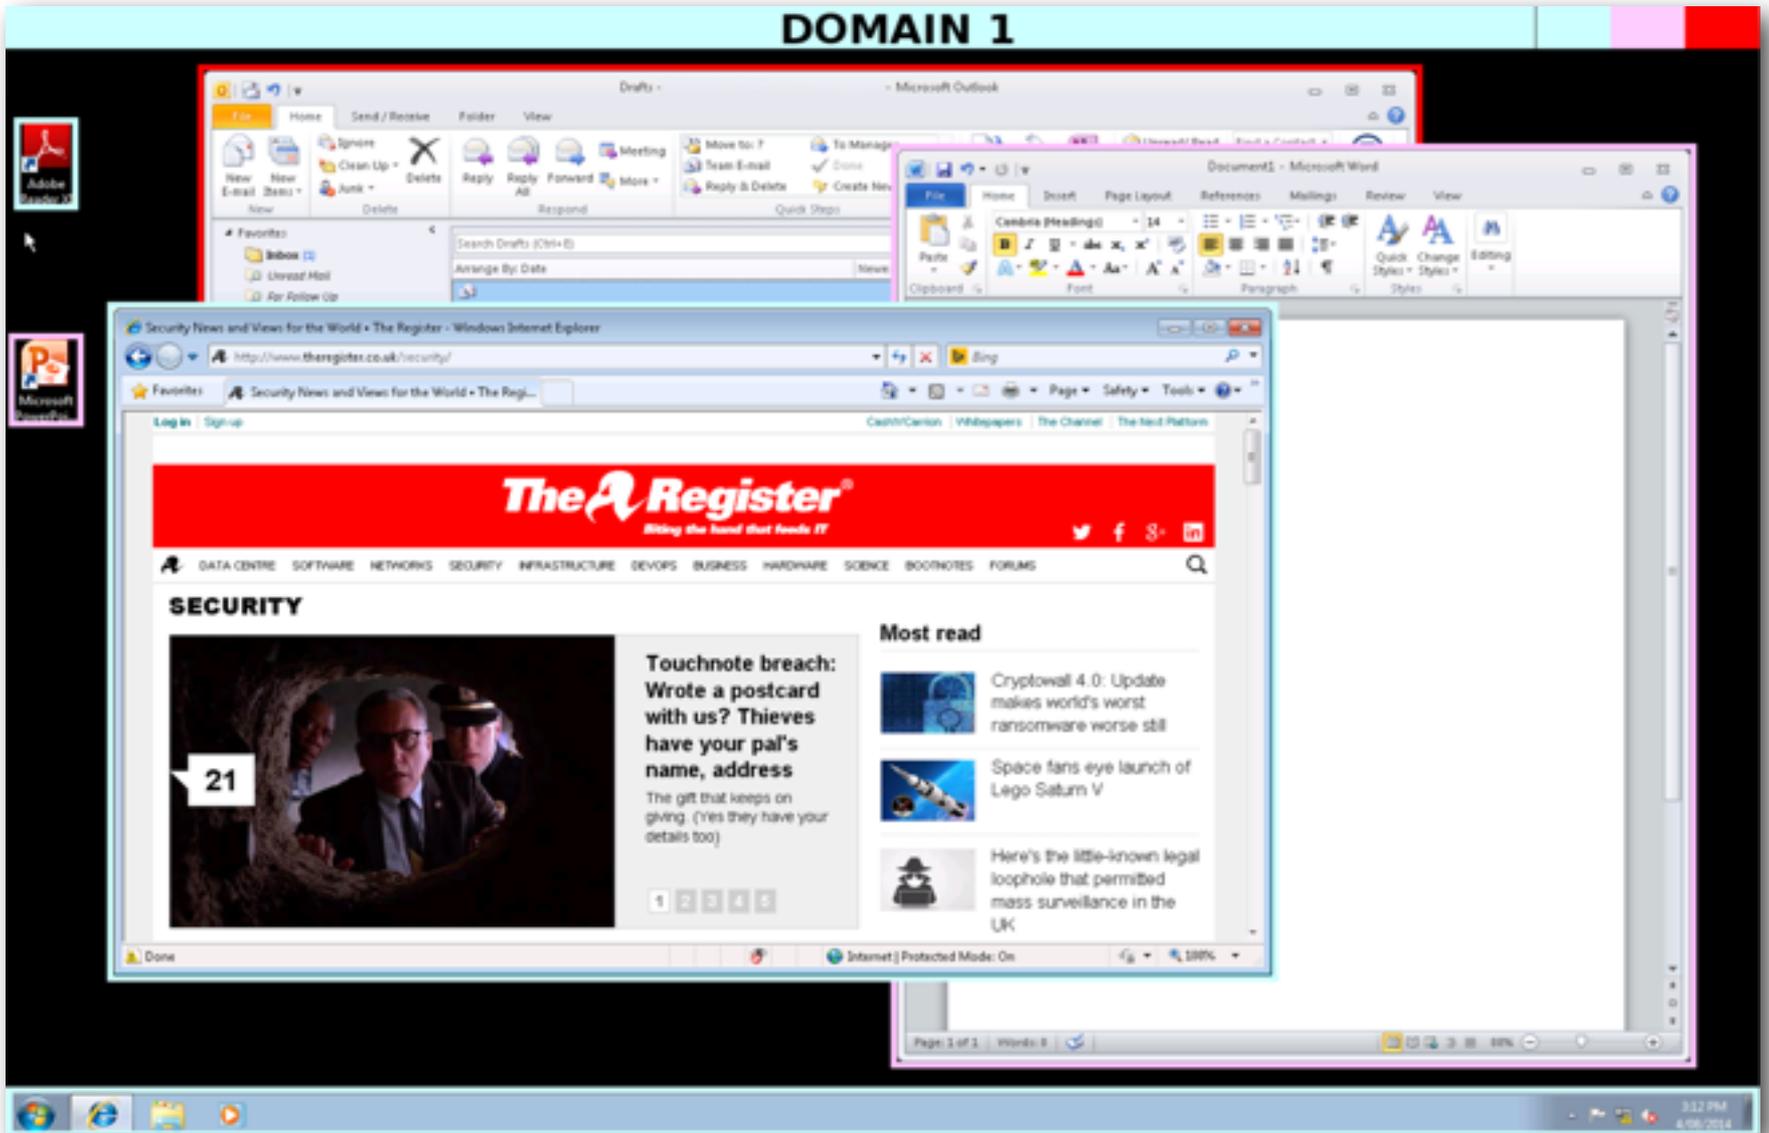


# CDDC

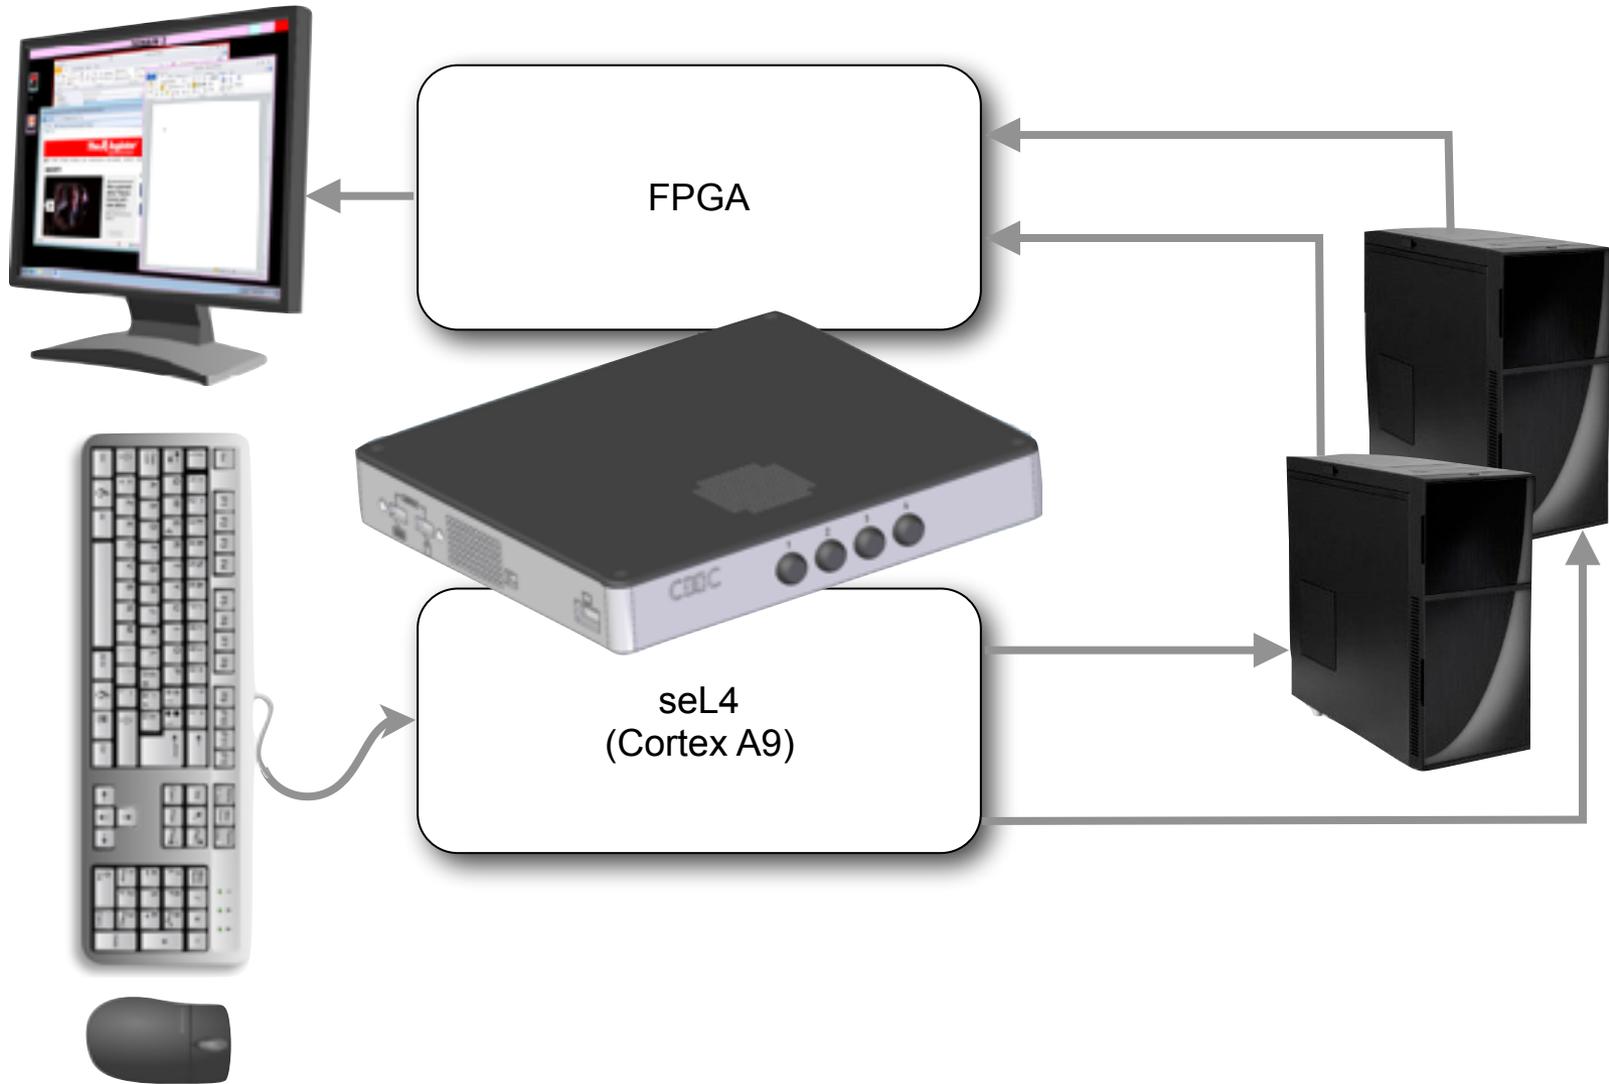


# CDDC User Experience

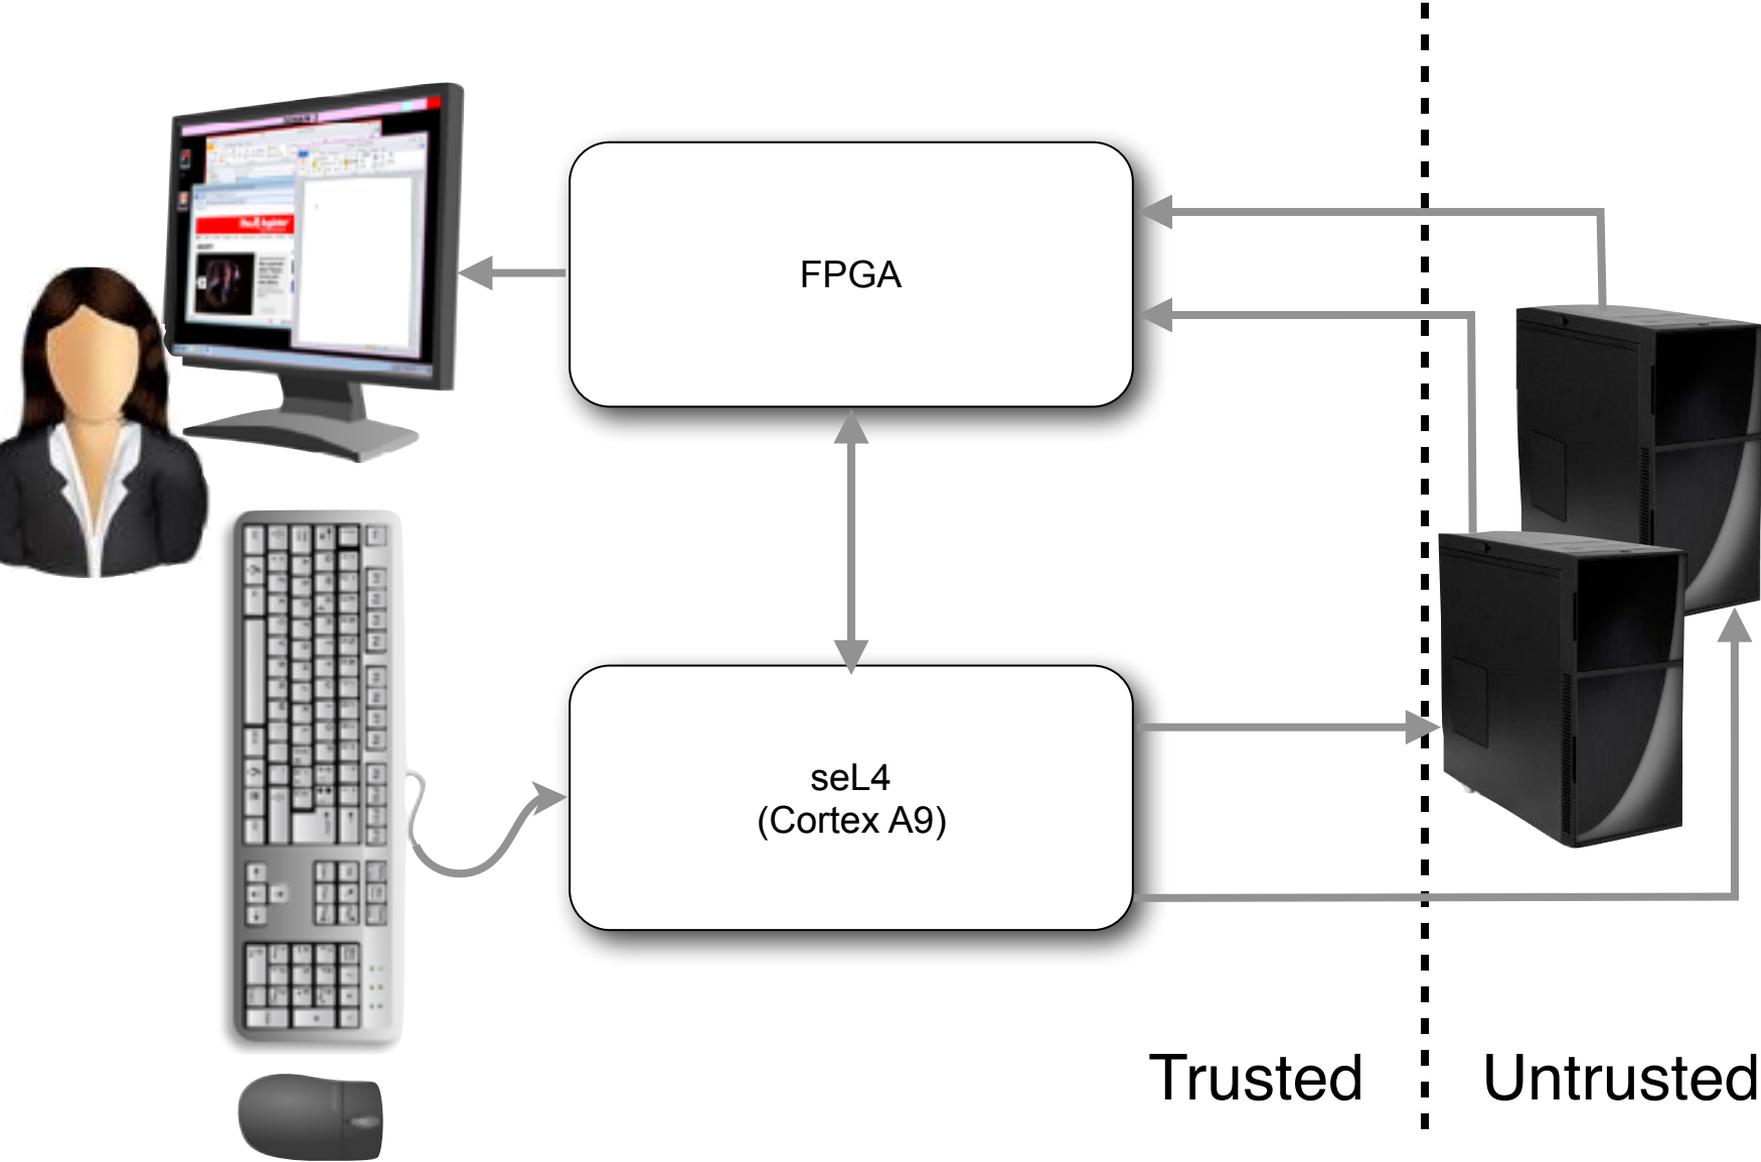
## DOMAIN 1



# CDDC Architecture

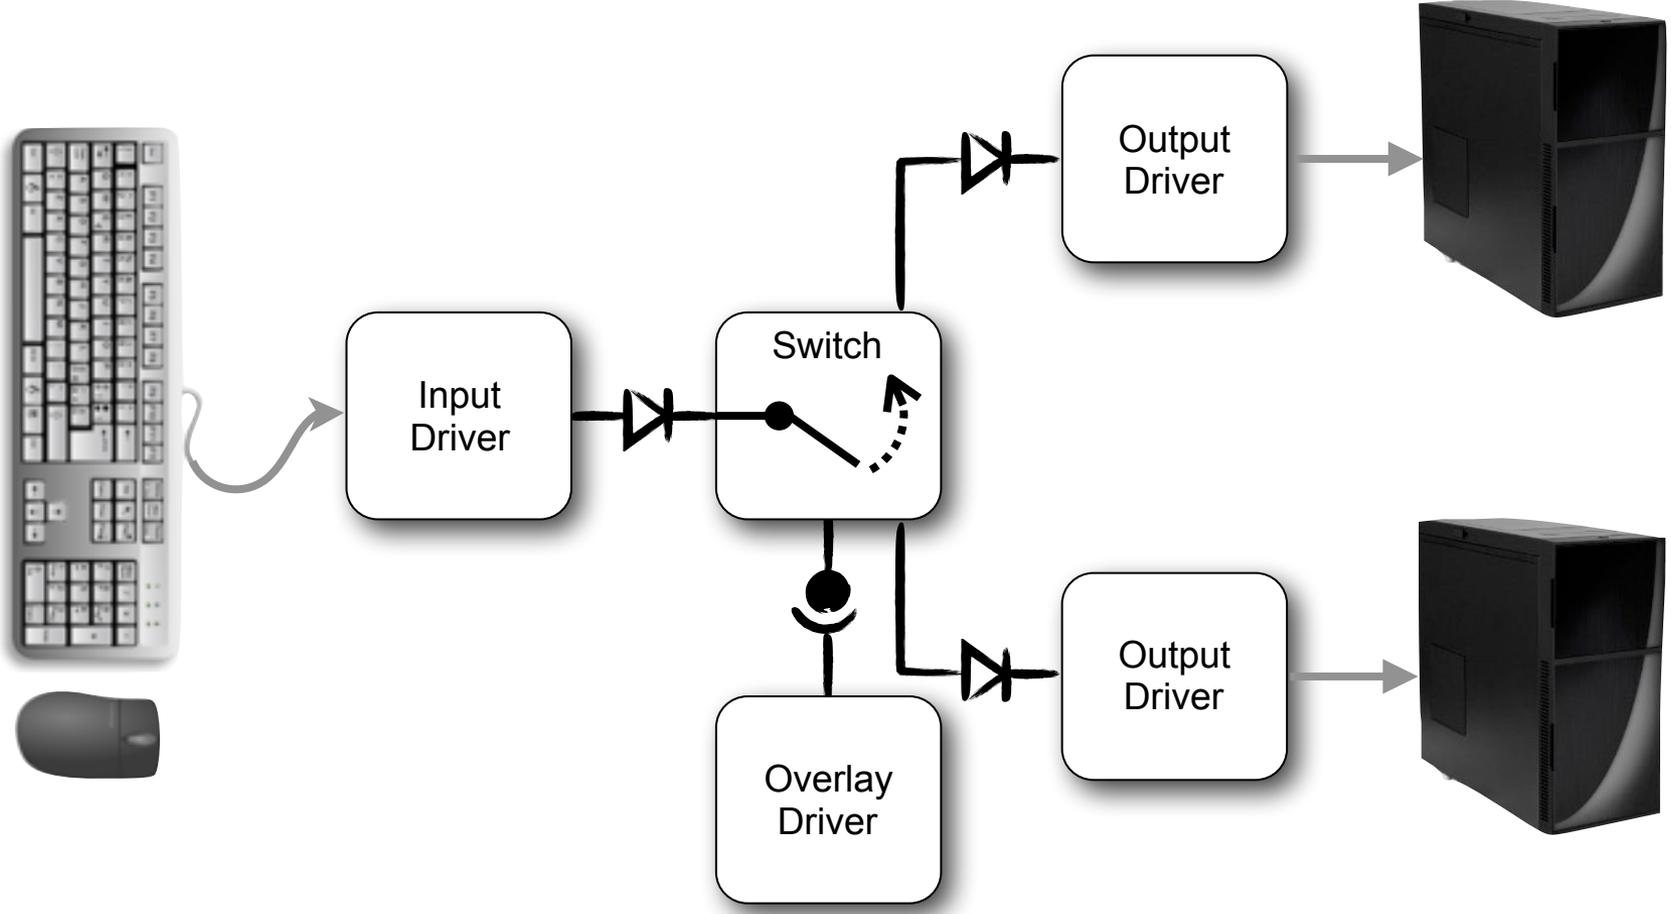


# Threat Model

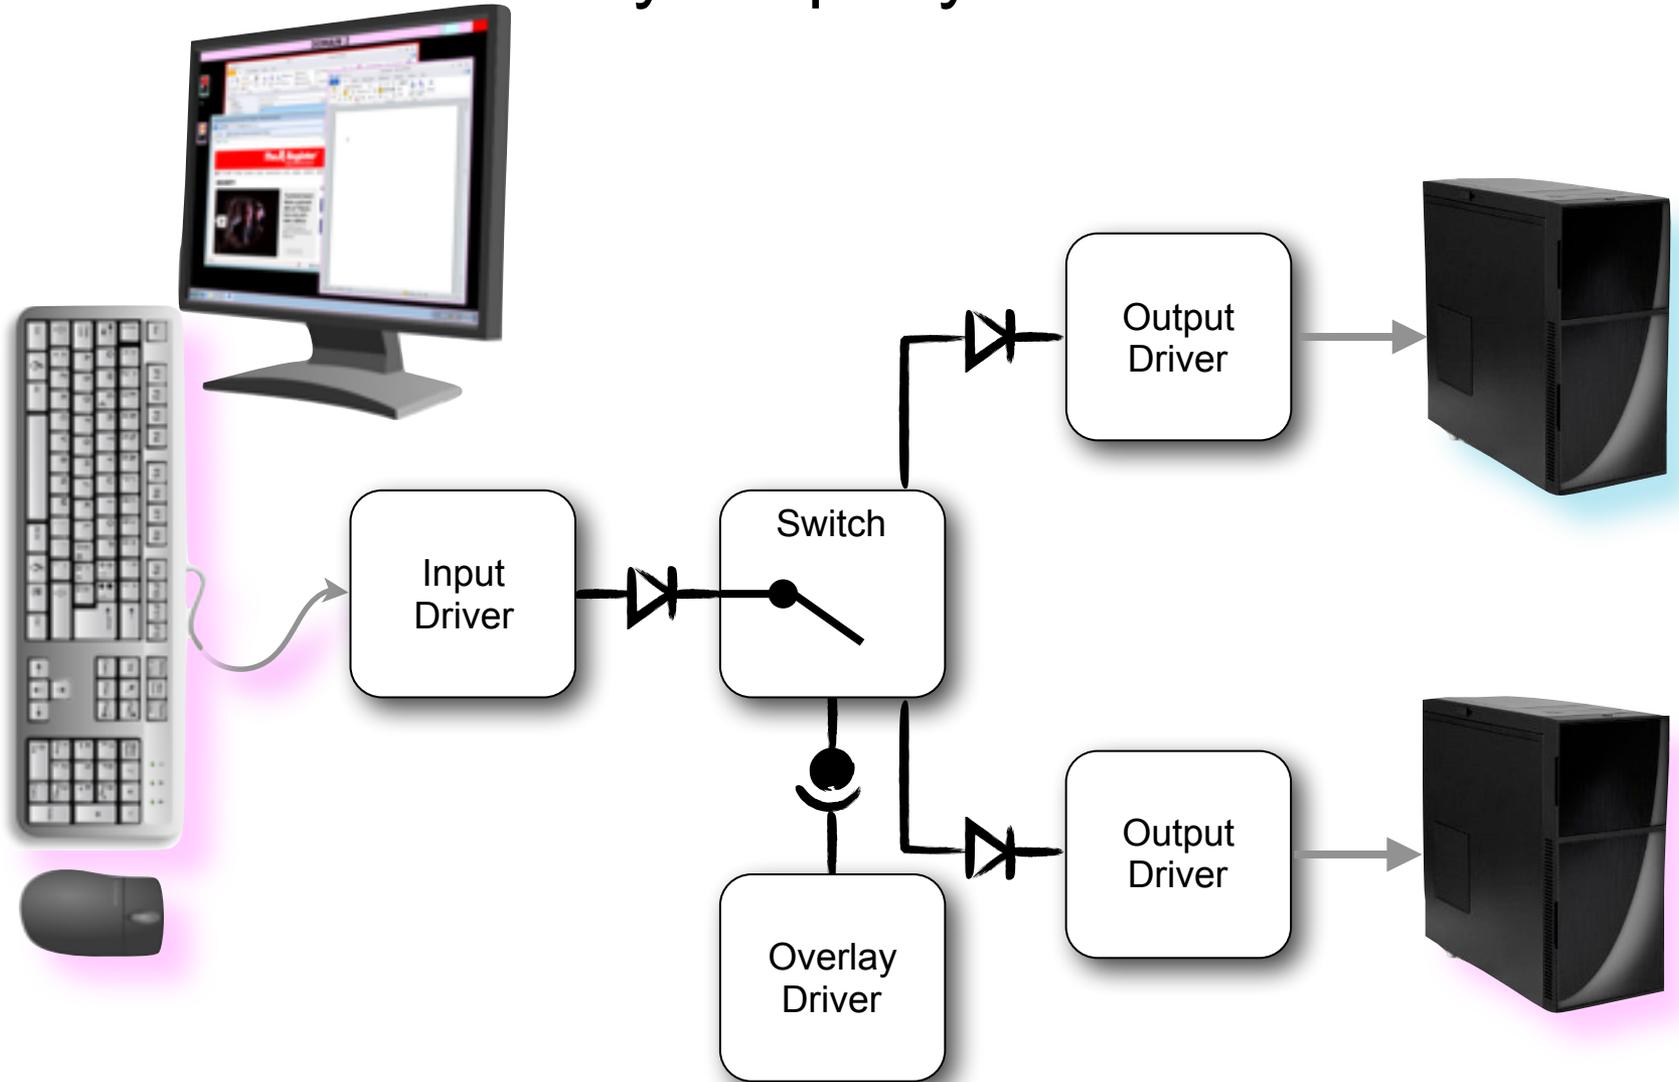


# seL4-Based Software Architecture

*(ignoring device administration and configuration, plus keyboard LED control)*

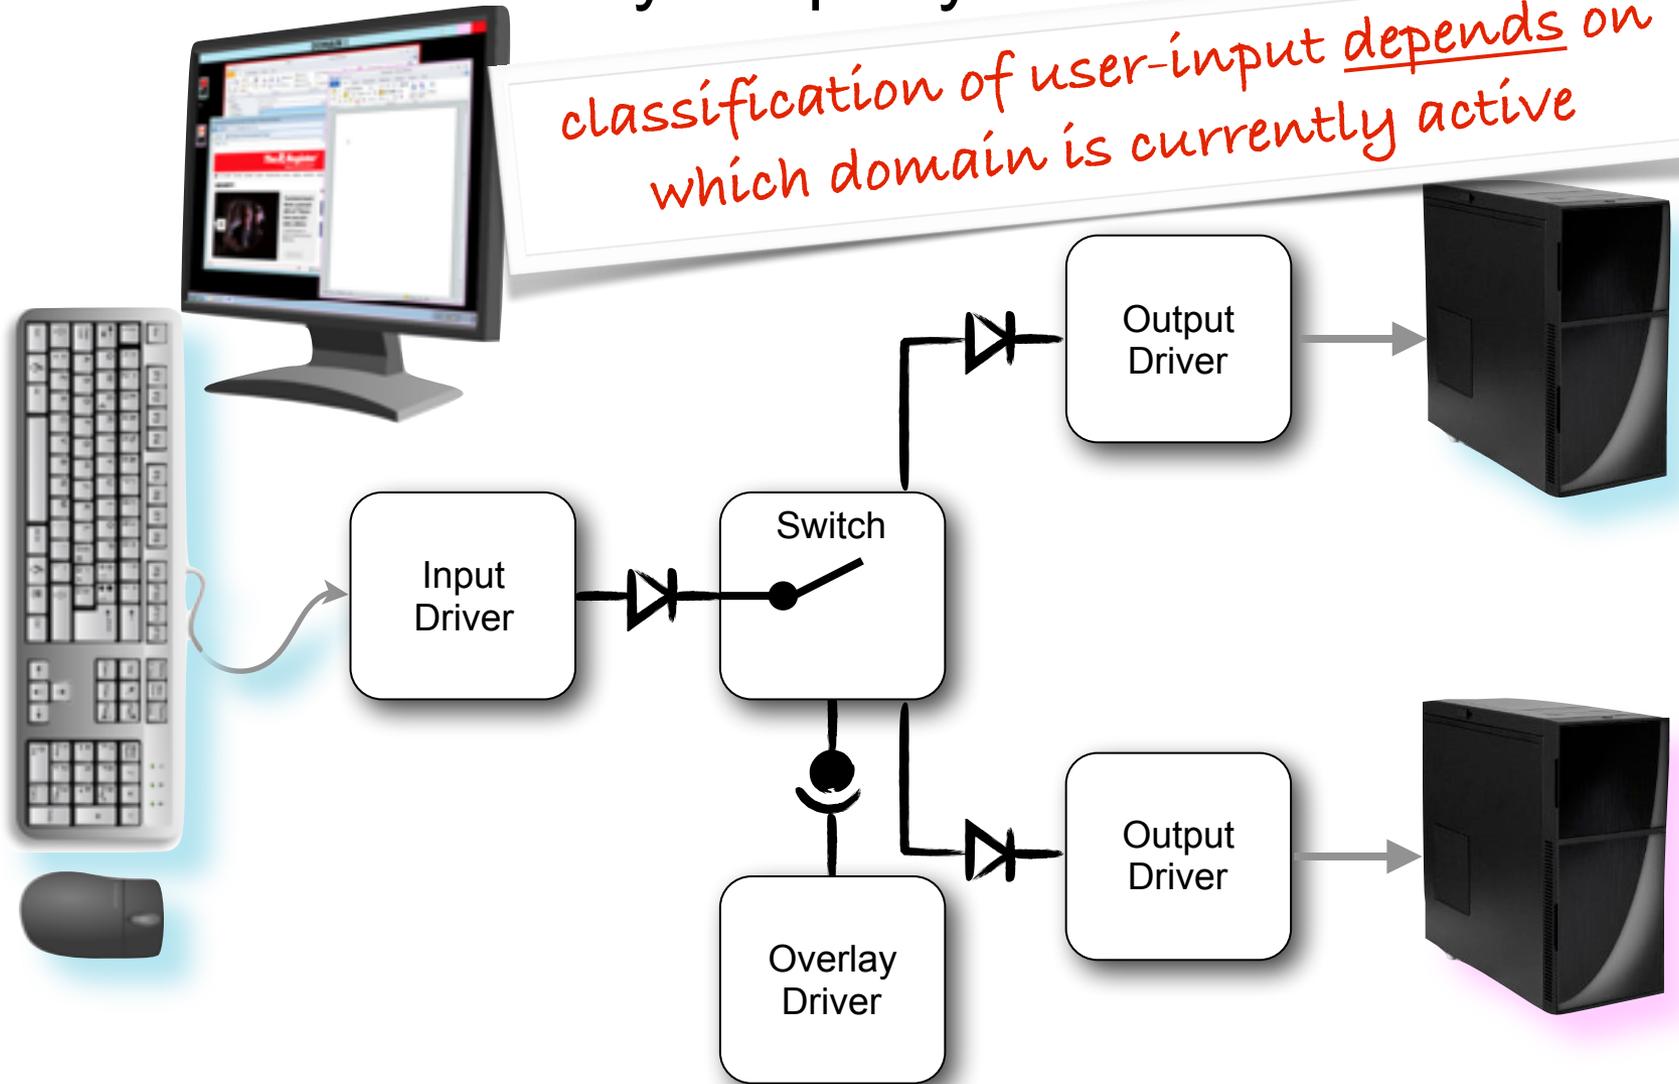


# Motivation: Security Property

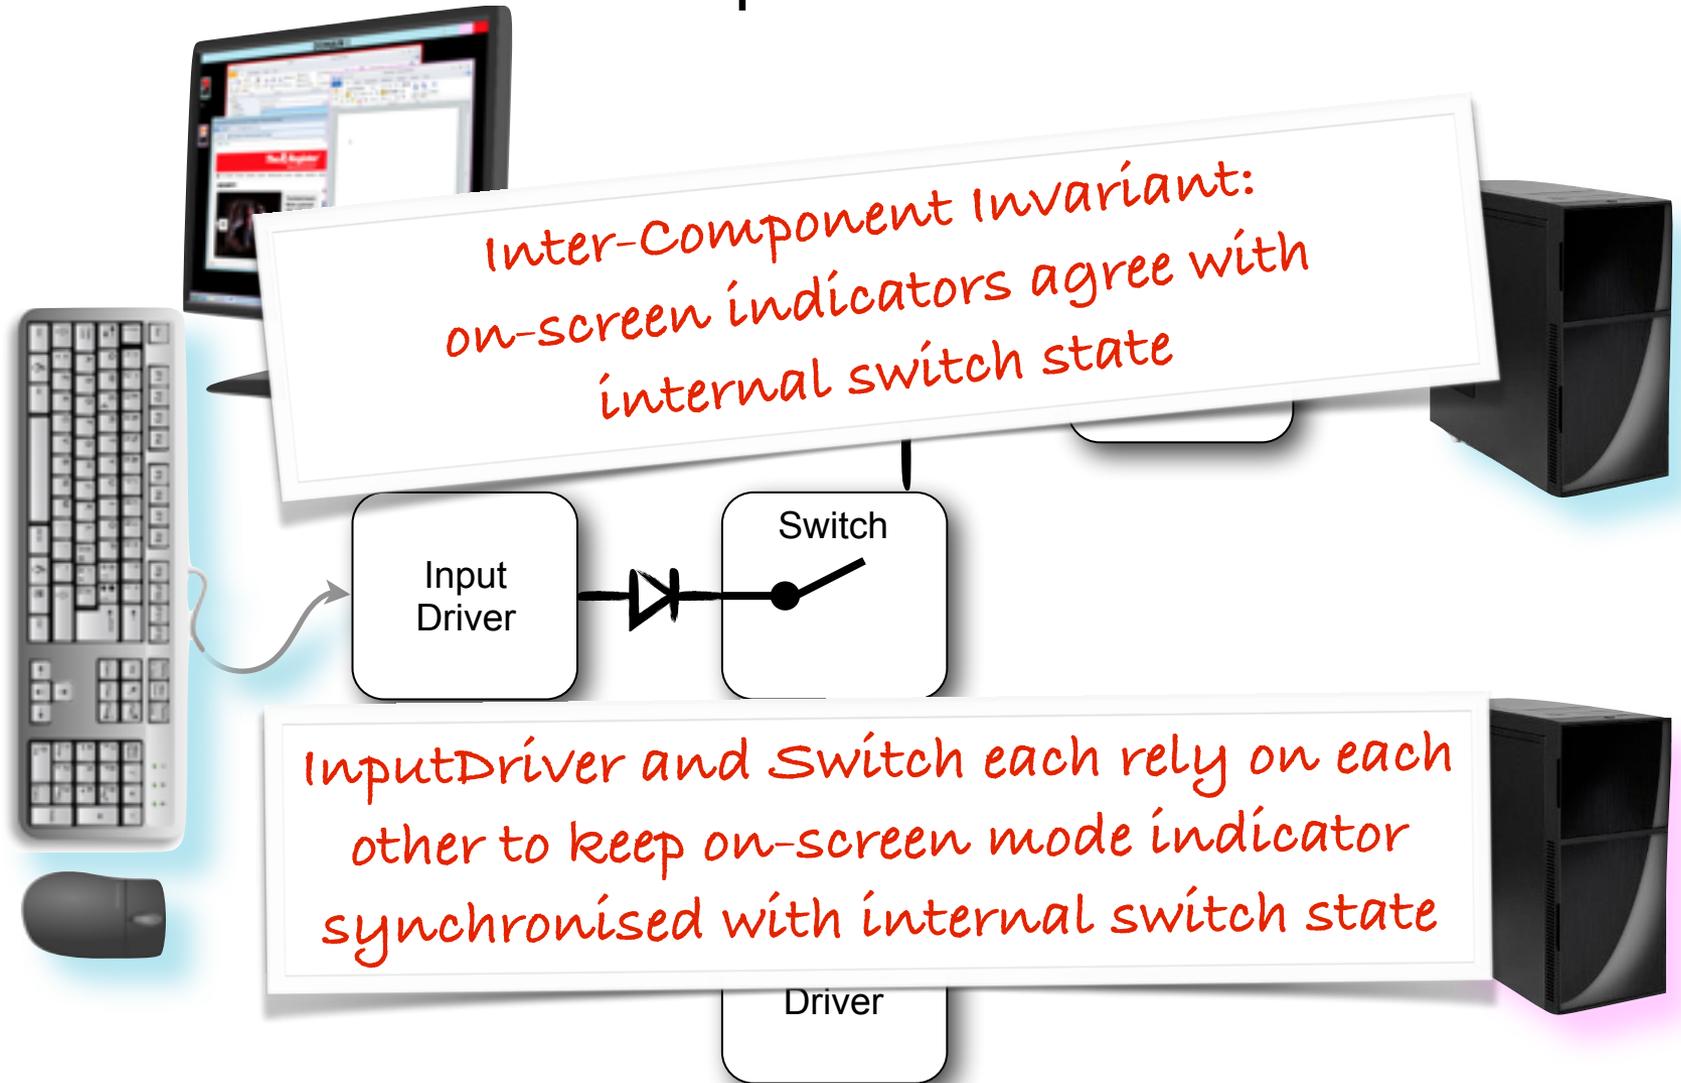


# Motivation: Security Property

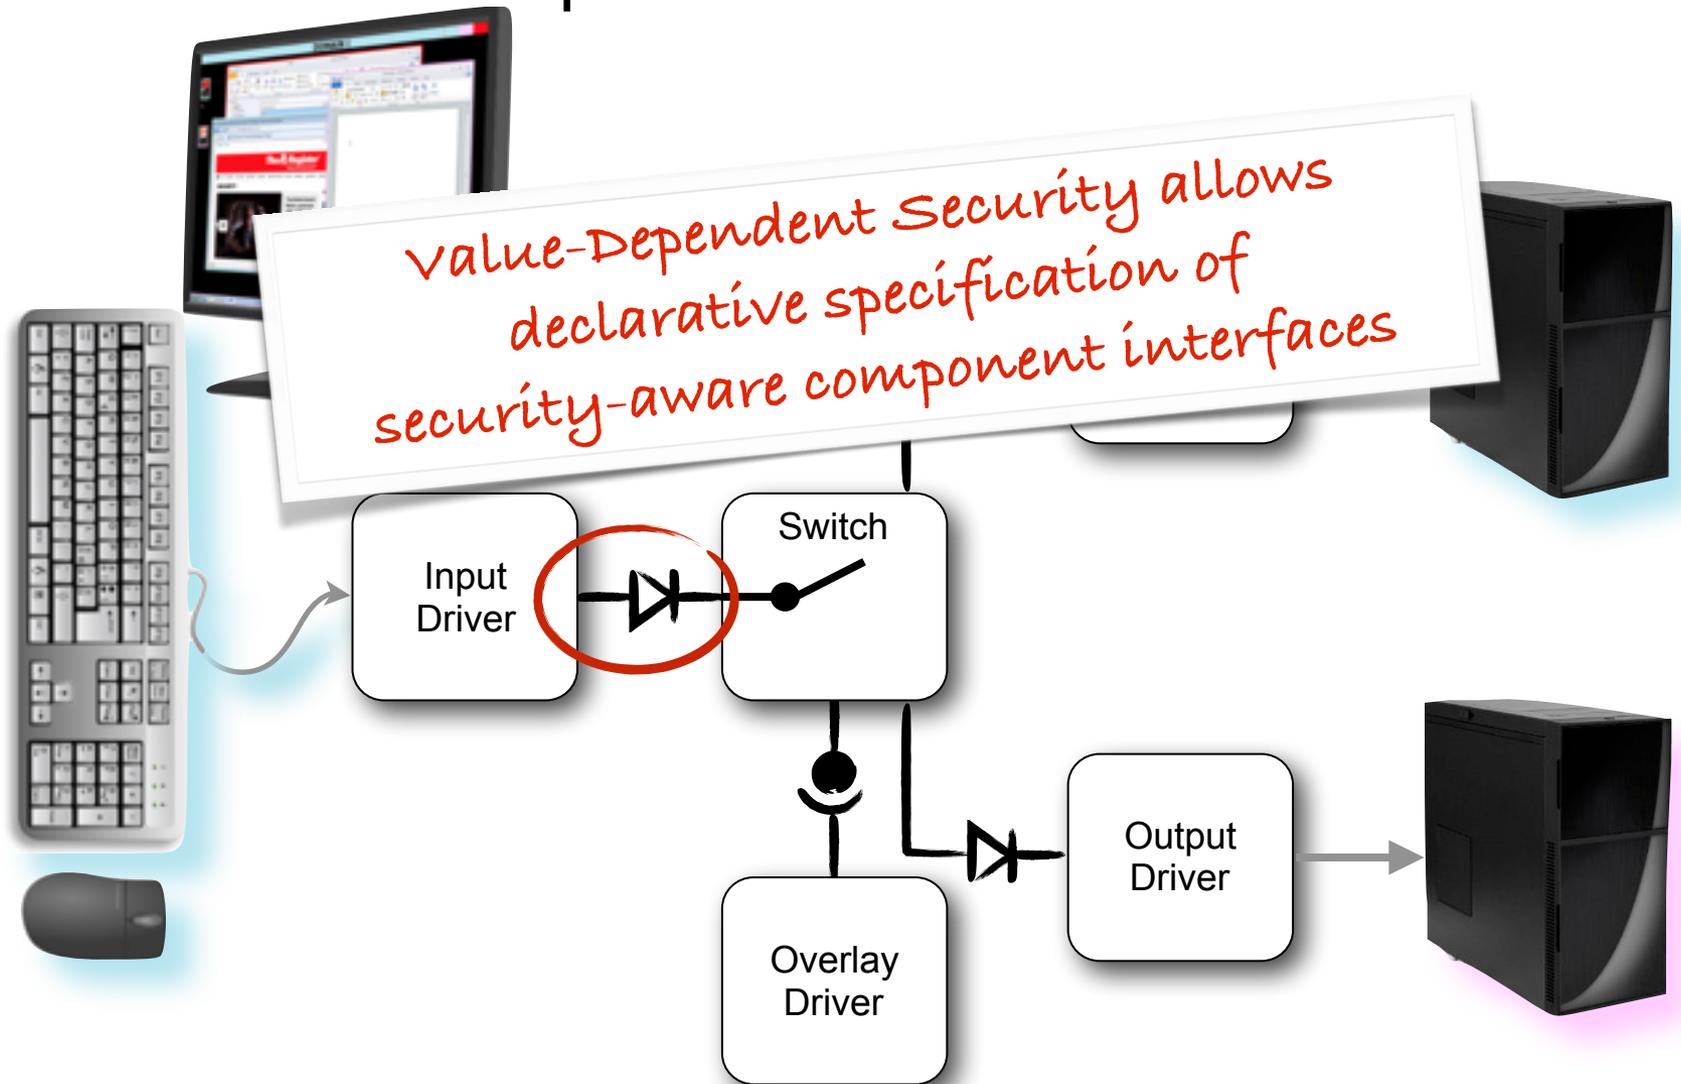
*classification of user-input depends on which domain is currently active*



# Motivation: Inter-Component Invariants



# Motivation: Component Interfaces



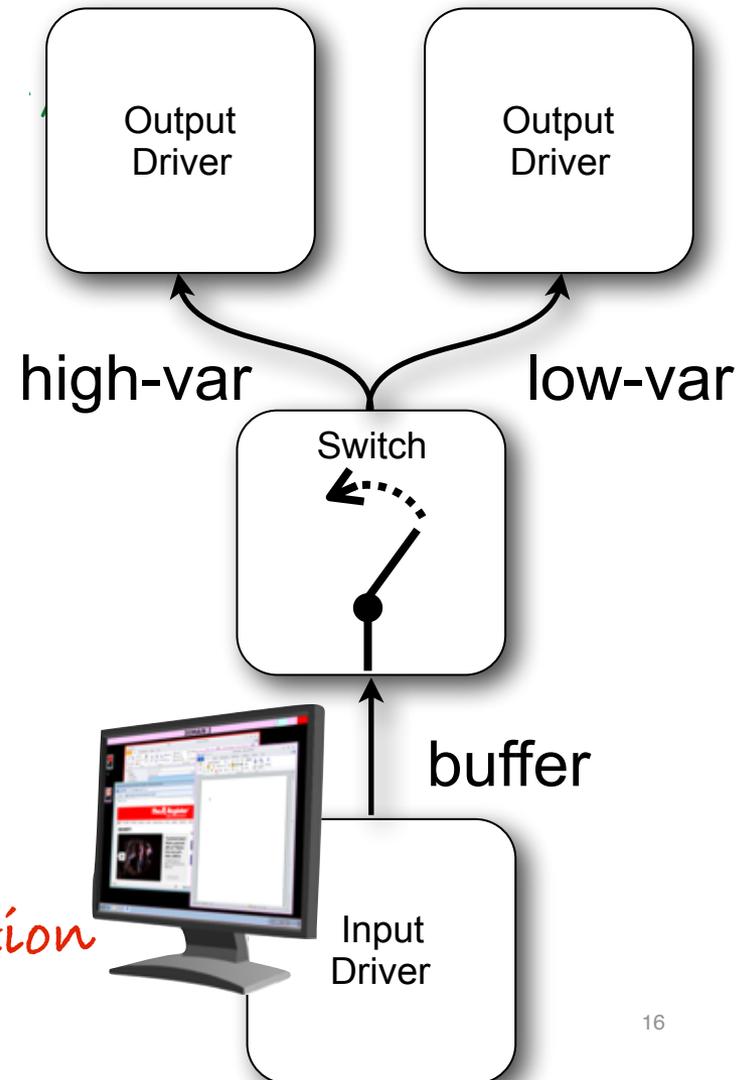
# Motivation: Shared Memory Concurrency

**Switch** (idealised, obviously):

```
temp := buffer;  
if sMode == 0 then  
  low-var := temp  
else  
  high-var := temp  
endif;  
temp := 0
```

sMode *is switch's mode*

dMode *defines buffer's classification*



# Compositional, Value-Dependent Security

*control variable (must be statically Low)*

```
var dMode // clas: Low
var high-var // clas: High
var low-var, sMode // clas: Low
var buffer // clas: dMode == 0 ? Low : High
lock l // NRW: {dMode, sMode, temp}
// inv: (dMode == sMode)
```

*value-dependent classification*

**lock(l):**

```
/* {dMode, sMode, temp} += AsmNoRW
   acquire (dMode == sMode) */
```

*lock footprint and invariant*

```
temp := buffer;
if sMode == 0 then
  low-var := temp
else
  high-var := temp
endif;
temp := 0
```

*(simplified, of course)*

*dynamic, local assumptions about environment, for compositional reasoning (acquired and released via locks)*

# THE COVERN LOGIC

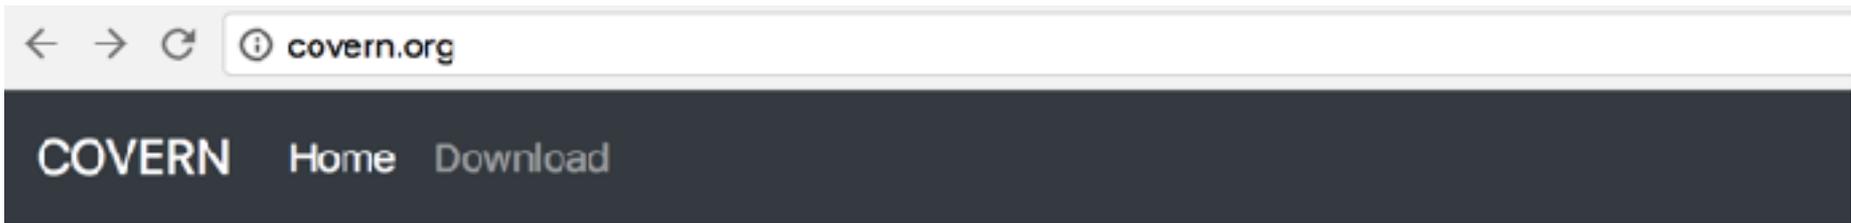


Robert Sison

$$\Gamma, L, P \{c\} \Gamma', L', P'$$


Kai Engelhardt

<http://covern.org>



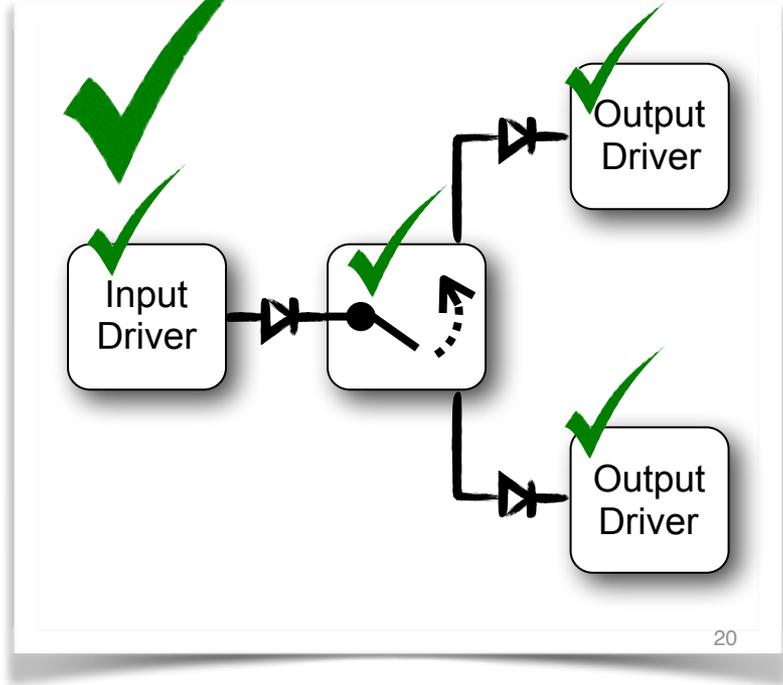
# The COVERN Project

Compositional Verification and Refinement of Noninterference

[More information »](#)

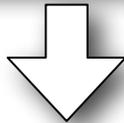
# Verification Framework: Requirements

Compositional Security Property



# Verification Framework: Requirements

Sound Program Security Logic



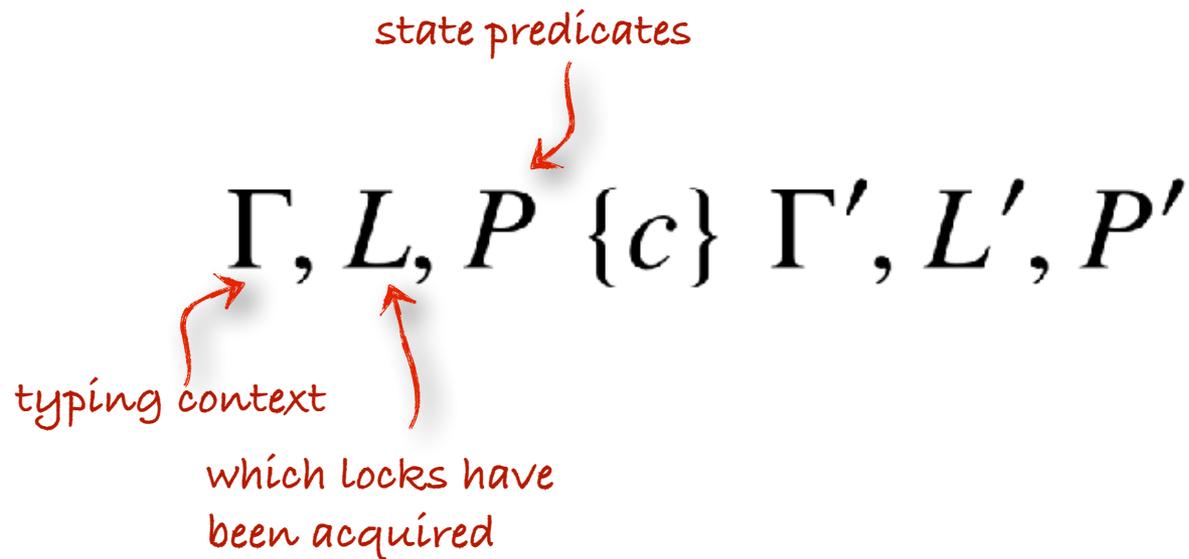
Compositional Security Property



```
lock(1);  
/* {dMode,sMode,temp} += AsmNoRW  
   acquire (dMode == sMode) */  
temp := buffer;  
if sMode == 0  
    low-var := t  
else  
    high-var := temp  
endif;  
temp := 0
```



# Types and Typing Judgements



**types depend on memory**

we encode them as state predicates

**type interpretation:**  $\llbracket t \rrbracket_{mem} \equiv \text{if } [t]_{mem} \text{ then Low else High}$

**subtype:**  $t \leq_{:P} t' \equiv (P \wedge t') \vdash t$

# Type System: Example in Action

```

[] {} true lock(1);

```

```

[t : true] {I} (sM = dM) temp := buffer;

```

```

[t : (dM = 0)] {I} (sM = dM) if sMode == 0 then

```

```

[t : (dM=0)] {I} (sM=dM ∧ sM=0) low-var := temp

```

$\Gamma$     $S$     $P$

```

else

```

```

    high-var := temp

```

```

endif;

```

```

temp := 0

```

Is (dM=0) a  
 subtype of **true**,  
 under  
 (sM=dM ∧ sM=0) ?

(sM=dM ∧ sM=0 ∧  
**true**) ⊢ (dM=0) ?

lock l protects sMode, dMode and temp, with invariant sMode = dMode

# CONCLUSION

General Logic for verifying IFC  
for Concurrent Programs

- concurrency  $\rightarrow$  compositionality  $\sim$   
timing-sensitivity, assume an annotated reasoning

**WRONG APPROACH?**

- inter-component contracts on shared memory:

- data invariants + value-dependent classifications

- concurrent IFC reasoning borrowing ideas from CSLs

- modelled and verified security-critical s/w functionality of CDC

Still lots more to be done.



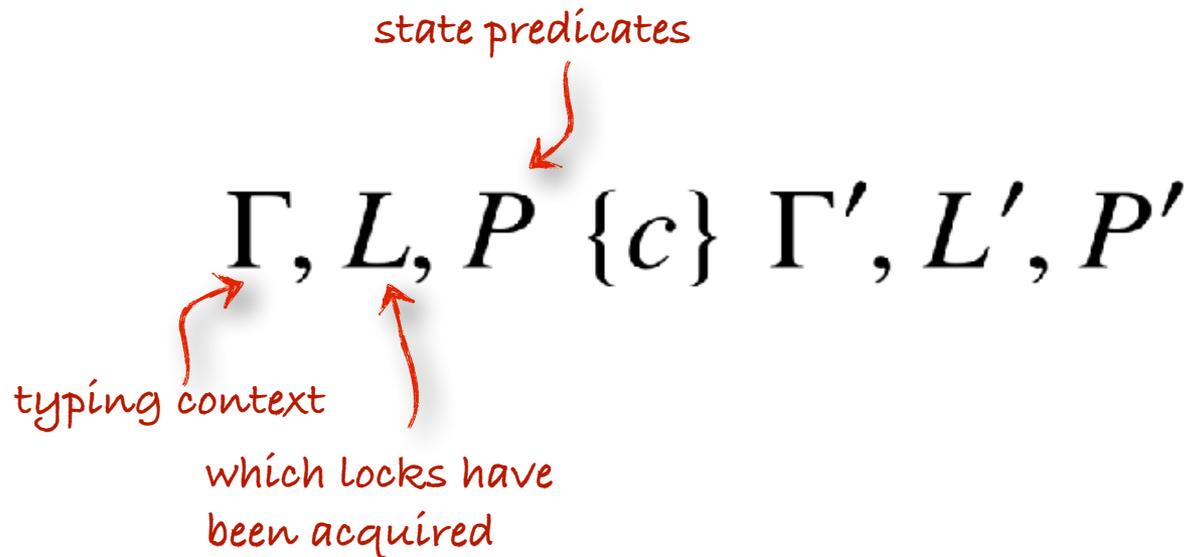
Daniel Schoepe



Andrei Sabelfeld

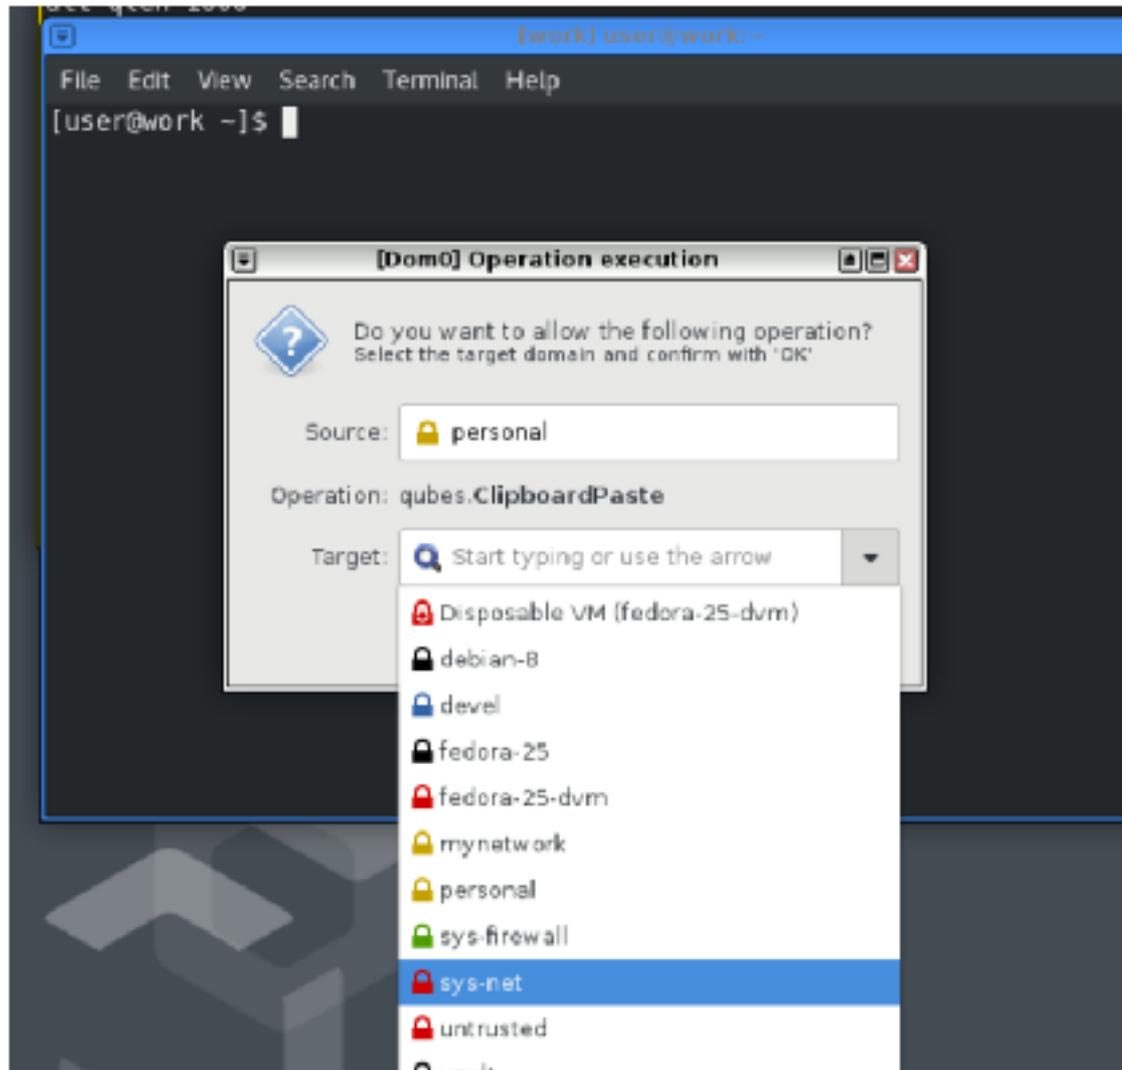
# Drawbacks

- No support for declassification
  - Compositional declassification in general not well understood



- Our logic does double duty:
  - Functional correctness (P and L)
  - Security ( $\Gamma$ )

# Declassification (in Qubes OS)



# Declassification

```
secret := ... // get secret from somewhere
output(H, "Declassify: " + secret);
answer := input(L);
if (answer = 1)
then declassify(secret);
```

We model allowed declassifications as *declassification predicates* on the current state  $s$  and the declassification event  $a$

- e.g.  $D(\alpha, s) \equiv last\_output(s, H) = \alpha \wedge last\_input(s, L) = 1$

- States  $s$  include full trace  $trace(s)$  of input/output events so far

# Defining Secure Declassification

There are so many declassification definitions.  
We use a modern, **knowledge-flavoured** definition.



Attacker can observe:

- initial Low (public) memory
- Low parts of the trace  $\tau$  (I/O on Low channels)

Attacker's **uncertainty**:

- we could have started in any state  $s$  that has the same initial Low memory and produced the same Low trace
- $uncertainty(s_0, t) \equiv$   
 $\{s_0' \mid s_0' =_L s_0 \wedge \exists t'. s_0' \rightarrow^* t' \wedge trace(t') =_L trace(t)\}$

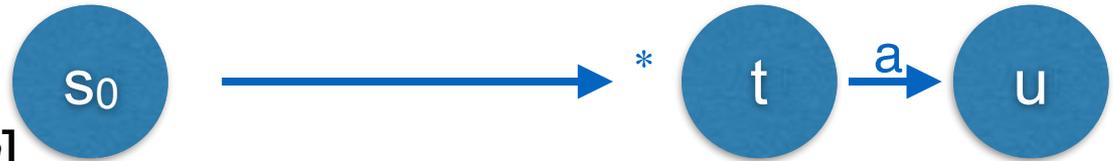
# Secure Declassification

**Security holds iff, whenever:**

- $s_0 \rightarrow^* t$

- $t \rightarrow u$

- $\text{trace}(u) = \text{trace}(t).[a]$



**Then:**

- If  $a$  is declassification then  $D(a,t)$  must hold
- Otherwise:  $\text{uncertainty}(s_0,t) \subseteq \text{uncertainty}(s_0,u)$  must hold

**Aside:** This is a bit weak for programs that branch on secrets:

```
if (birthYear > 2000)
```

```
  Declassify(birthDay);
```

```
else
```

```
  Declassify(birthMonth);
```

**DISALLOWED**

# How to prove security compositionally?

A compositional knowledge-based definition?

Too hard

**Observation:** Type systems / program logic just defines and proves a (relational) invariant.

**Why not just do that directly?**

**Also, let's decouple the functional correctness and security parts.**

# Example

```
Thread 1:  
acquire(lock);  
if (src = 0)  
then buffer := input(L)  
else buffer := input(H)  
release(lock);
```

# Annotations (1)

```
Thread 1:  
acquire(lock);  
if (src = 0)  
then buffer := input(L)  
else buffer := input(H)  
release(lock);
```

```
Thread 2:  
acquire(lock);  
if (dest = 0)  
then output(L, buffer)  
else output(H, buffer)  
release(lock);
```

*dest = src*

*src = 0*

```
Thread 3:  
acquire(lock);  
src = (src = 0) ? 1 : 0;  
dest = src;  
release(lock);
```

*src = 0 ⇒ buffer = input(L) ∧  
src = 1 ⇒ buffer = input(H)*

# Annotations (2)

Thread 1:  
acquire(lock);  
if (src = 0)  
then buffer := input(L)  
else buffer := input(H)  
release(lock);

*has\_lock(lock)*

Thread 2:  
acquire(lock);  
if (dest = 0)  
then output(L, buffer)  
else output(H, buffer)  
release(lock);

*has\_lock(lock)*

Thread 3:  
acquire(lock);  
src = (src = 0) ? 1 : 0;  
dest = src;  
release(lock);

*has\_lock(lock)*

$src = 0 \Rightarrow buffer = input(L) \wedge$   
 $src = 1 \Rightarrow buffer = input(H)$

*dest = src*

*src = 0*

# Annotations (3)

Thread 1:

```
acquire(lock);  
if (src = 0)  
then buffer := input(L)  
else buffer := input(H)  
release(lock);
```

*has\_lock(lock)*

Thread 2:

```
acquire(lock);  
if (dest = 0)  
then output(L, buffer)  
else output(H, buffer)  
release(lock);
```

*has\_lock(lock)*

*dest = src*

*src = 0*

Thread 3:

```
acquire(lock);  
src = (src = 0) ? 1 : 0;  
dest = src;  
release(lock);
```

*has\_lock(lock)*

*dest = src*

$src = 0 \Rightarrow buffer = input(L) \wedge$   
 $src = 1 \Rightarrow buffer = input(H)$

# Annotated Programs

Commands  $c$  carry precondition annotations  $A$

$c ::=$  **skip**  
|  $c_1 ; c_2$   
|  $\{A\} x := e$   
|  $\{A\} x :=_D e$   
|  $\{A\} \mathbf{out}(l, e)$   
|  $\{A\} x \leftarrow l$   
|  $\{A\} \mathbf{if} (e) \mathbf{then} c_1 \mathbf{else} c_2$   
|  $\{A\} \mathbf{while} e \{l\} \mathbf{do} c$   
|  $\{A\} \mathbf{acquire}(lck)$   
|  $\{A\} \mathbf{release}(lck)$

# Type System for Annotated Programs

**Intuition:** Precise annotations allow dependency (data flow) information to be inferred

T-OUT

$$\frac{\forall s, s'. s =_{\ell} s' \wedge s \in A \wedge s' \in A \Rightarrow \llbracket e \rrbracket(s) = \llbracket e \rrbracket(s')}{\vdash \{A\} \text{out}(\ell, e)}$$

T-DECL

$$\frac{\forall s. s \in A \Rightarrow \mathcal{D}(\llbracket e \rrbracket(s), s)}{\vdash \{A\} \text{declassify}(e)}$$

# Soundness Theorem

local reasoning use your favourite concurrent program verification method

→ If every thread is well-typed,  
and all annotations hold when the threads are run in parallel,  
→ Then their parallel composition satisfies the security property

(so far we have used Owicki-Gries (Prensa Nieto '02) and Rely-Guarantee)

(generate annotations using your favourite program analysis methods, since generation doesn't need to be trusted anyway)

# Type System Example

---

⊢ **if** (*dest* = 0) **then** **out**(*L*, *buffer*) **else** **out**(*H*, *buffer*)

# Typing Rules

$$\begin{array}{c}
 \text{T-SKIP} \\
 \hline
 \ell \vdash \mathbf{skip}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{T-LASSIGN} \\
 x \in \text{dom}(\mathcal{L}) \quad \mathcal{L}(x) \sqsubseteq \mathcal{L}(e) \\
 \hline
 \ell \vdash \{A\} x = e
 \end{array}
 \qquad
 \begin{array}{c}
 \text{T-UASSIGN} \\
 x \notin \text{dom}(\mathcal{L}) \\
 \hline
 \ell \vdash \{A\} x = e
 \end{array}$$
  

$$\begin{array}{c}
 \text{T-DECL} \\
 \forall g. A(g) \Rightarrow \mathcal{D}(\mathcal{L}(e), \mathcal{L}(x), \llbracket e \rrbracket(g), \{A\} x =_D e) \\
 \hline
 \ell \vdash \{A\} x =_D e
 \end{array}
 \qquad
 \begin{array}{c}
 \text{T-OUT} \\
 A \models e \\
 \hline
 \ell \vdash \{A\} \mathbf{out}(\ell', e)
 \end{array}$$
  

$$\begin{array}{c}
 \text{T-LIN} \\
 x \in \text{dom}(\mathcal{L}) \quad \mathcal{L}(x) \sqsubseteq \ell' \\
 \hline
 \ell \vdash \{A\} x \leftarrow \ell'
 \end{array}
 \qquad
 \begin{array}{c}
 \text{T-ULIN} \\
 x \notin \text{dom}(\mathcal{L}) \\
 \hline
 \ell \vdash \{A\} x \leftarrow \ell'
 \end{array}$$
  

$$\begin{array}{c}
 \text{T-IF} \\
 A \models e \quad \ell \vdash c_1 \quad \ell \vdash c_2 \\
 \hline
 \ell \vdash \{A\} \mathbf{if} (e) \mathbf{then} c_1 \mathbf{else} c_2
 \end{array}
 \qquad
 \begin{array}{c}
 \text{T-WHILE} \\
 A \models e \quad I \models e \quad \ell \vdash c \\
 \hline
 \ell \vdash \{A, I\} \mathbf{while} e \mathbf{do} c
 \end{array}$$

where:

$$A \models e = \forall m, m'. m =_{\ell} m' \wedge m \in A \wedge m' \in A \rightarrow \llbracket e \rrbracket(m) = \llbracket e \rrbracket(m')$$

# c.f. Murray et al., EuroS&P 2018

$$\frac{\text{modifiable}_L(x) \quad \text{readable}_L(e) \quad \Gamma \vdash e : t \quad \text{stable}_L(t) \quad P' = \text{post}_L(x := e, P) \quad x \notin \text{NRW}(L) \longrightarrow t \leq_{P'} \mathcal{L}(x)}{\Gamma, L, P \{x := e\} \Gamma[x \mapsto t], L, P'} \text{ASSIGN}_1$$

$$\frac{x \in C \quad \text{modifiable}_L(x) \quad \text{readable}_L(e) \quad \Gamma \vdash e : t \quad t \leq_P \text{LOW} \quad (\forall v \in \text{dom}(\Gamma). x \notin \text{vars}(\Gamma(v))) \quad P' = \text{post}_L(x := e, P) \quad \text{secure\_update}}{\Gamma, L, P \{x := e\} \Gamma', L, P'} \text{ASSIGN}_C$$

$$\frac{\text{readable}_L(e) \quad \Gamma \vdash e : t \quad t \leq_P \text{LOW} \quad \Gamma, L, (P \wedge e) \{c_1\} \Gamma_1, L', P_1 \quad \Gamma, L, (P \wedge \neg e) \{c_2\} \Gamma_2, L', P_2 \quad \Gamma_1 =_{P_1} \Gamma' \quad \Gamma_2 =_{P_2} \Gamma' \quad P_1 \vdash P' \quad P_2 \vdash P' \quad (\Gamma_1, L', P_1) \leq (\Gamma', L', P') \quad (\Gamma_2, L', P_2) \leq (\Gamma', L', P')}{\Gamma, L, P \{\text{if } e \text{ then } c_1 \text{ else } c_2\} \Gamma', L', P'} \text{IFLOW}$$

$$\frac{\text{readable}_L(e) \quad \Gamma \vdash e : t \quad t \leq_P \text{LOW} \quad \Gamma, L, (P \wedge e) \{c\} \Gamma, L, P}{\Gamma, L, P \{\text{while } e \text{ do } c\} \Gamma, L, P} \text{WHILE}$$

$$\frac{\Gamma, L, P \{c_1\} \Gamma', L', P' \quad \Gamma', L', P' \{c_2\} \Gamma'', L'', P''}{\Gamma, L, P \{c_1 ; c_2\} \Gamma'', L'', P''} \text{SEQ} \qquad \frac{}{\Gamma, L, P \{\text{skip}\} \Gamma, L, P} \text{SKIP}$$

$$\frac{\Gamma' = \Gamma \ominus \ell}{\Gamma, L, P \{\text{lock}(\ell)\} \Gamma', L \cup \{\ell\}, (P \wedge \text{lockinv}(\ell))} \text{LOCKACQ}$$

$$\frac{\ell \in L \quad P \vdash \text{lockinv}(\ell) \quad \Gamma' = \Gamma \ominus \ell \quad P' = P \ominus \ell \quad \text{side\_condition}}{\Gamma, L, P \{\text{unlock}(\ell)\} \Gamma', L - \{\ell\}, P'} \text{LOCKREL}$$

# Isabelle Formalisation

<b>Theory</b>	<b>LoC</b>
Language	1k
Security defs.	150
Type system & soundness	2k
Owicki-Gries	600
Rely-Guarantee	1.5k
Delimited Release	600
Misc	500
<b>Total</b>	<b>6.5k</b>

c.f. COVERN is 10K

More modular, extensible; integrates with existing O-G framework; easy to automate type checking via Eisbach

# Future

- **Extending to programs that branch on secrets**
  - How far can we get with Hoare logic annotations here?
- **Extending to weak memory**
  - Feasible now that functional correctness has been decoupled
    - Since “morally” IFC type system doesn’t care about the order in which statements are executed
  - Use your favourite weak memory verification techniques...
- **Larger case studies**
  - e.g. redo the CDDC verification in the new framework

# CONCLUSION

Decouple functional correctness and security for verifying IFC for concurrent programs

- supports compositional declassification
- inter-component contracts on shared memory:
  - expressed as Hoare logic assertions
- use your favourite concurrent program techniques
- 

Still lots more to be done.

# Thank You



[toby.murray@unimelb.edu.au](mailto:toby.murray@unimelb.edu.au)



<http://people.eng.unimelb.edu.au/tobym>



[@tobycmurray](https://twitter.com/tobycmurray)

# Connection to Delimited Release

- Delimited release [Sabelfeld & Myers 2003] ensures only information intended to be declassified can be leaked
- Can be soundly embedded into our model with side condition that *occurrence* of declassification events does not depend on secrets.
- Makes sense for programs that do not branch on secrets
- But unclear how to extend to programs that do branch on secrets