

# Dynamically Provisioning Isolation in Hierarchical Architectures

Kevin Falzon  
kevin.falzon@ec-spride.de

Eric Bodden  
eric.bodden@ec-spride.de

September 9, 2015



EC SPRIDE  
EUROPEAN CENTER FOR  
SECURITY AND PRIVACY BY DESIGN



TECHNISCHE  
UNIVERSITÄT  
DARMSTADT

EUROPEAN CENTER FOR SECURITY AND PRIVACY BY DESIGN

## Side channels




## Covert channels




---


<https://www.flickr.com/photos/inkvision/2868388856>

<http://jeffberezny.com/wp-content/uploads/2014/02/passing-paper-notes-in-class.jpg>







Core-level attacks  
(AES key inference, ElGamal path detection)




Machine-level attacks  
(memory bus covert channel, branch predictor, /proc monitoring)



Cross-VM attacks  
(VCPU sharing, scheduler attacks, core-level attacks)



Hypervisor-level attacks  
(fingerprinting, memory deduplication, machine-level attacks)



Network-level attacks  
(OSI layer hiding, remote timing attacks)



## Observation

Channels are scoped.



Isolate entities at a granularity matching a perceived threat

- Isolate **on demand** rather than statically
- Finer granularity → better utilisation

Isolate entities at a granularity matching a perceived threat

- Isolate **on demand** rather than statically
- Finer granularity → better utilisation
- Note: focus is on *hard isolation*

Look into migration at several system levels:

- Move **processes** between CPUs, vCPUs and OS environments
- Move **virtual machines** between physical machines

Enhance approach via additional technologies:

- Hardware event counters
- Post-copy migration
- Process containers

Developed the SAFEHAVEN framework

- Networks of *agent* and *probe* processes
  - Agents control migrations
  - Probes serve as event sources
- Interface with **cpusets**, **criu** and **libvirt**
- Originally in C, now in Erlang



Erlang

Developed the SAFEHAVEN framework

- Networks of *agent* and *probe* processes

*Agents* control migrations

*Probes* serve as event sources

- Interface with **cpuset**, **criu** and **libvirt**
- Originally in C, now in Erlang



Two case studies:

- Machine-wide covert channel by Wu et al.
- Multi-level moving target defense


Developed the SAFEHAVEN framework

- Networks of *agent* and *probe* processes
  - *Agents* control migrations
  - *Probes* serve as event sources
- Interface with **cpuset**, **criu** and **libvirt**
- Originally in C, now in Erlang



Two case studies:

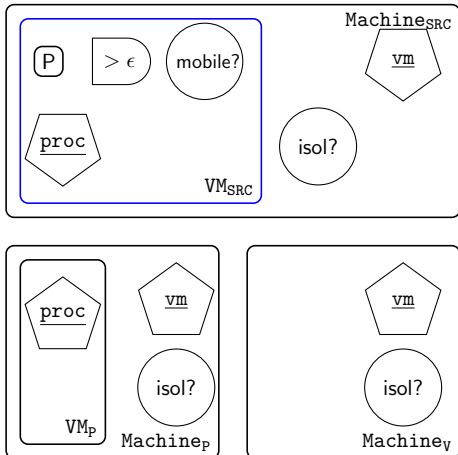
- Machine-wide covert channel by Wu et al.
- Multi-level moving target defense

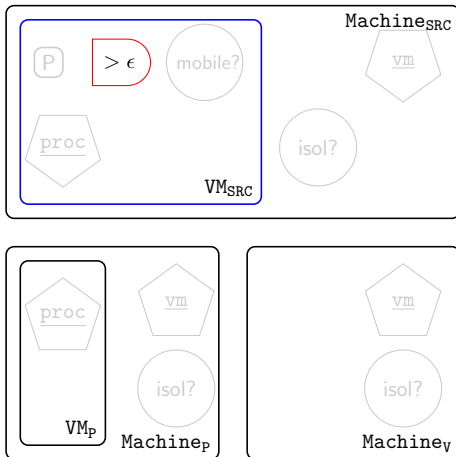


Process A attempts to communicate with B via a covert channel

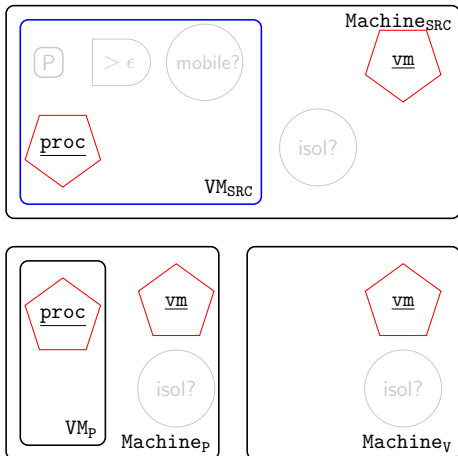
- Misaligned atomic accesses causes bus lock
- Modulate access times



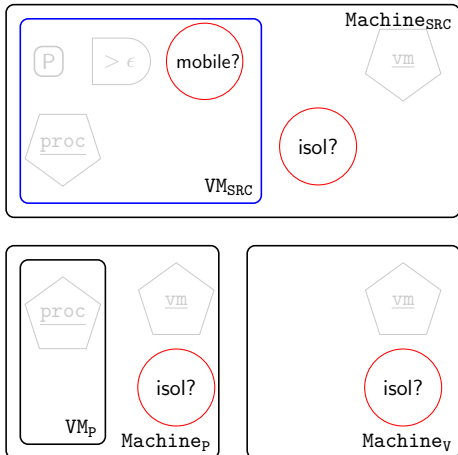




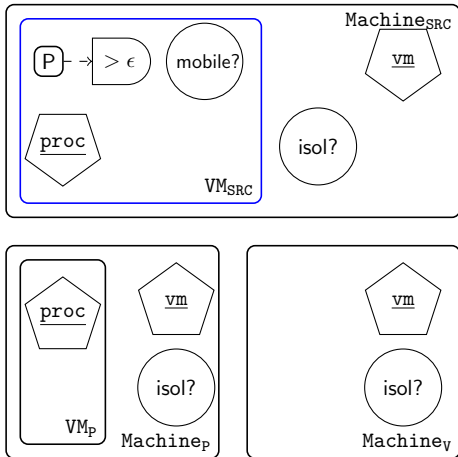
Probe (Hardware counter)



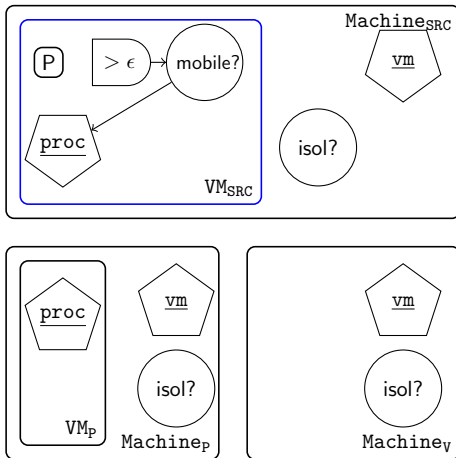
Agents



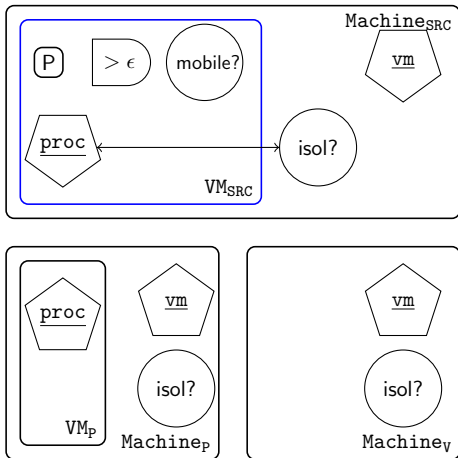
Probes/Predicates



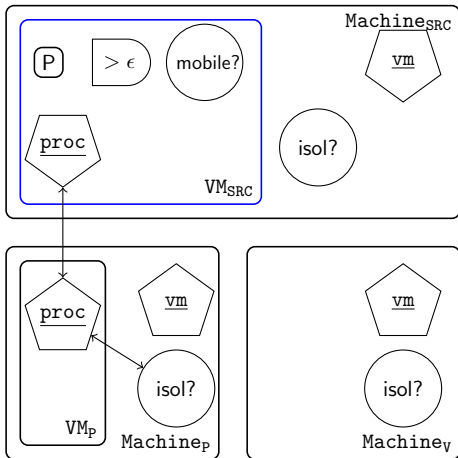
Process emits events at rate exceeding  $\epsilon$



If process can be moved ...




Check isolation criteria

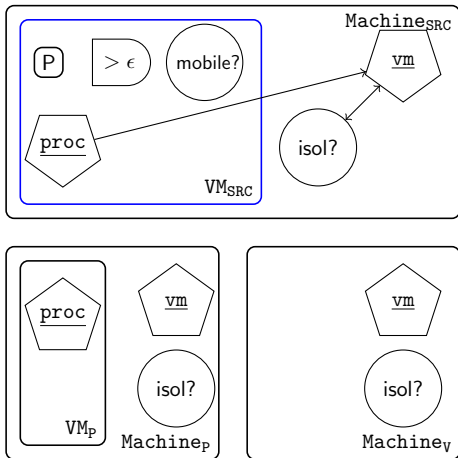


Find target system matching criteria






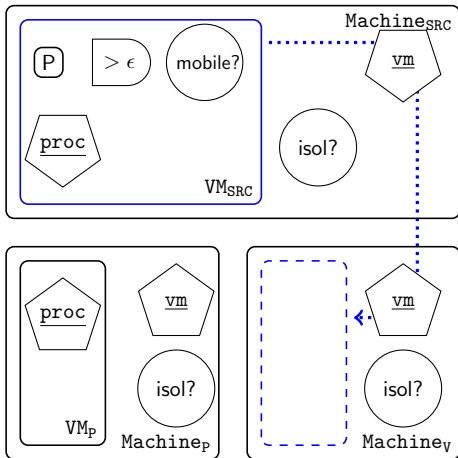
Migrate process



Process cannot move; resolve at VM level



Consult VM-level agents to identify isolated target




Perform VM migration

Observe frequency of misaligned atomic memory operations using hardware counters.

$$\text{Sample Time} = \frac{\text{Period}}{\# \text{ Processes}}$$


- Fixed period (1000, 2500ms)
- 120-200 processes, low duty cycle
- Detects vast majority in first pass




Observe frequency of misaligned atomic memory operations using hardware counters.

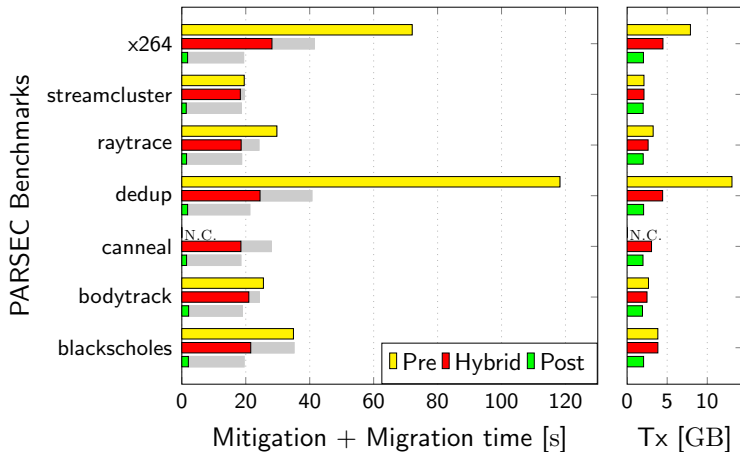
$$\text{Sample Time} = \frac{\text{Period}}{\# \text{ Processes}}$$

- Fixed period (1000, 2500ms)
- 120-200 processes, low duty cycle\*
- Detects vast majority in first pass



\*can be even faster/more accurate!







	Min	Max	Geo. Mean	Arith. Mean
Detection, period = 1s	0.015	3.16	0.54	0.72
Post-copy migration	1.281	2.13	1.47	1.48
Total (seconds)	<b>1.296</b>	<b>5.29</b>	<b>2.01</b>	<b>2.20</b>

	Min	Max	Geo. Mean	Arith. Mean
Detection, period = 1s	0.015	3.16	0.54	0.72
Post-copy migration*	1.281	2.13	1.47	1.48
Total (seconds)	<b>1.296</b>	<b>5.29</b>	<b>2.01</b>	<b>2.20</b>

\*can be even faster!

## Condensed results:

- Migrating processes to cores/vCPUs takes  $\approx 25\text{ms}$
- Container checkpoint and restore times are workload-dependent – order of seconds
  - ▶ Checkpoint:  $\approx 4\text{s} - 24\text{s}$ , mean  $\approx 10\text{s}$
  - ▶ Restore:  $\approx 1\text{s} - 8\text{s}$ , mean  $\approx 2\text{s}$
  - ▶ Memory pressure largest C/R performance factor
  - ▶ `rsync` incurred massive performance penalty ( $\approx 40\text{s}$ )
- Post-copy VM migration times are very consistent
  - ▶  $\approx 20\text{s}$  total transfer for a 2GB VM
  - ▶ Migrating every 30s halved performance; 60s  $\approx 30\%$  overhead

## Conclusions

1. Conjunction of mechanisms makes the approach viable. Other factors include total system capacity and economic concerns.
2. Post-copy makes it possible to mitigate attacks at the coarsest level of granularity in a timely manner.
3. Process migration is theoretically the most general mitigation, but it is not yet sufficiently robust.

## Future work

- Automatic property generation
- Finding intermediate levels of granularity

Thank you!