Confidential Containers and the Pitfalls of Runtime Attestation

Tobin Feldman-Fitzthum

IBM Research

Decoupled measurement

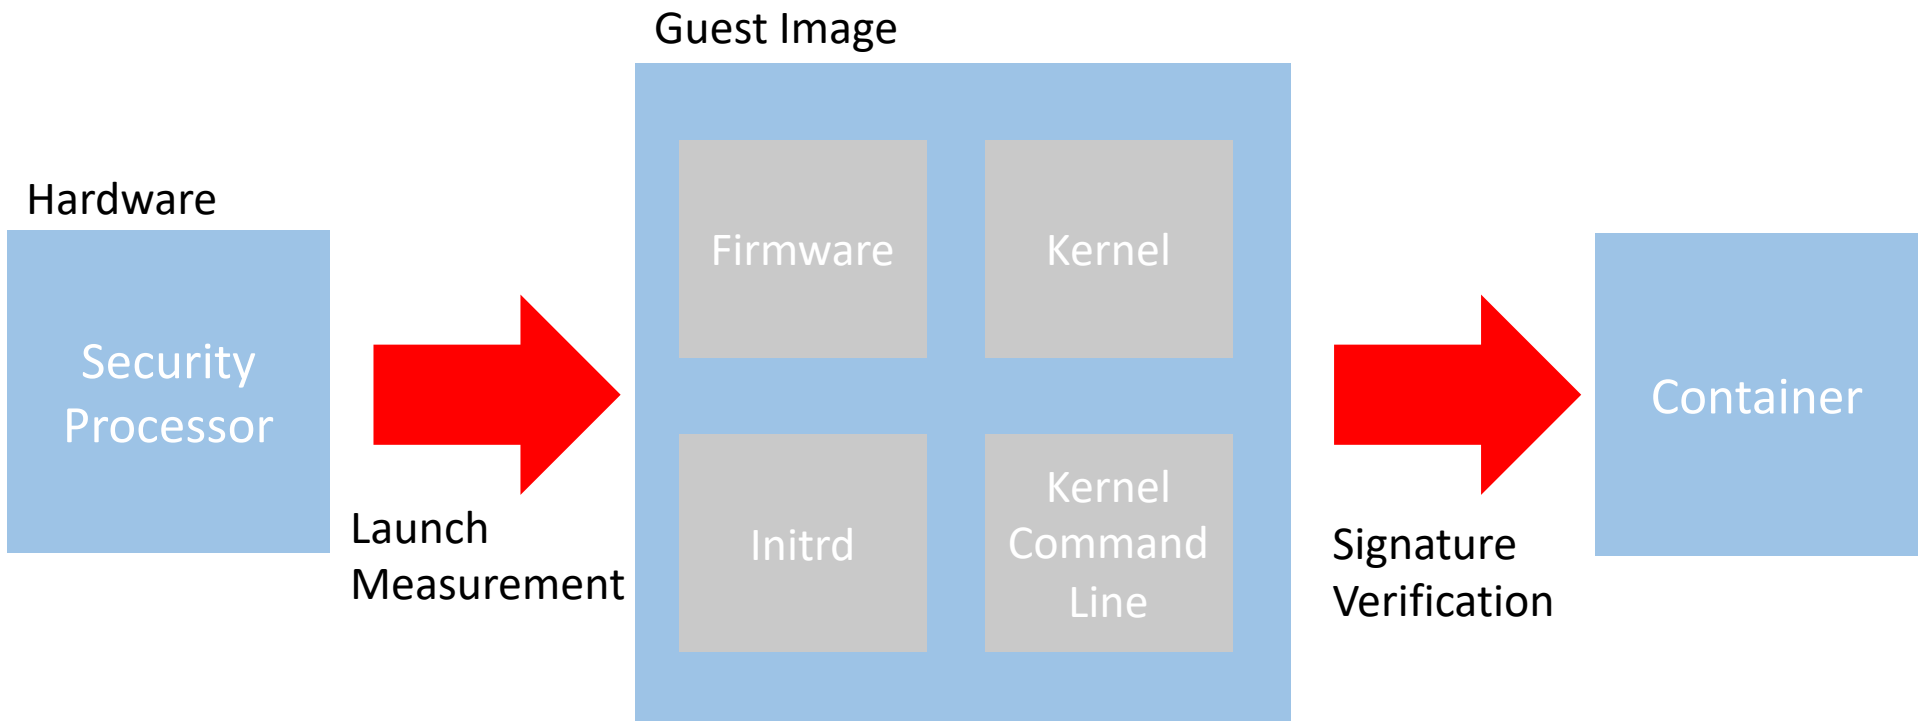

Evidence Factory Attacks

- Attestation Reports can’t be tampered with
- But what if you could make your own
  - Could you impersonate a valid guest and get secrets?
  - Just start your own VM
  - But the launch measurement would be wrong
  - Or, if the launch measurement is correct, the VM is not malicious

- Can a malicious VM manufacture a valid attestation report?
  - With a valid launch measurement with the correct fw, initrd, kernel, cmdline

Evidence Factory Attacks

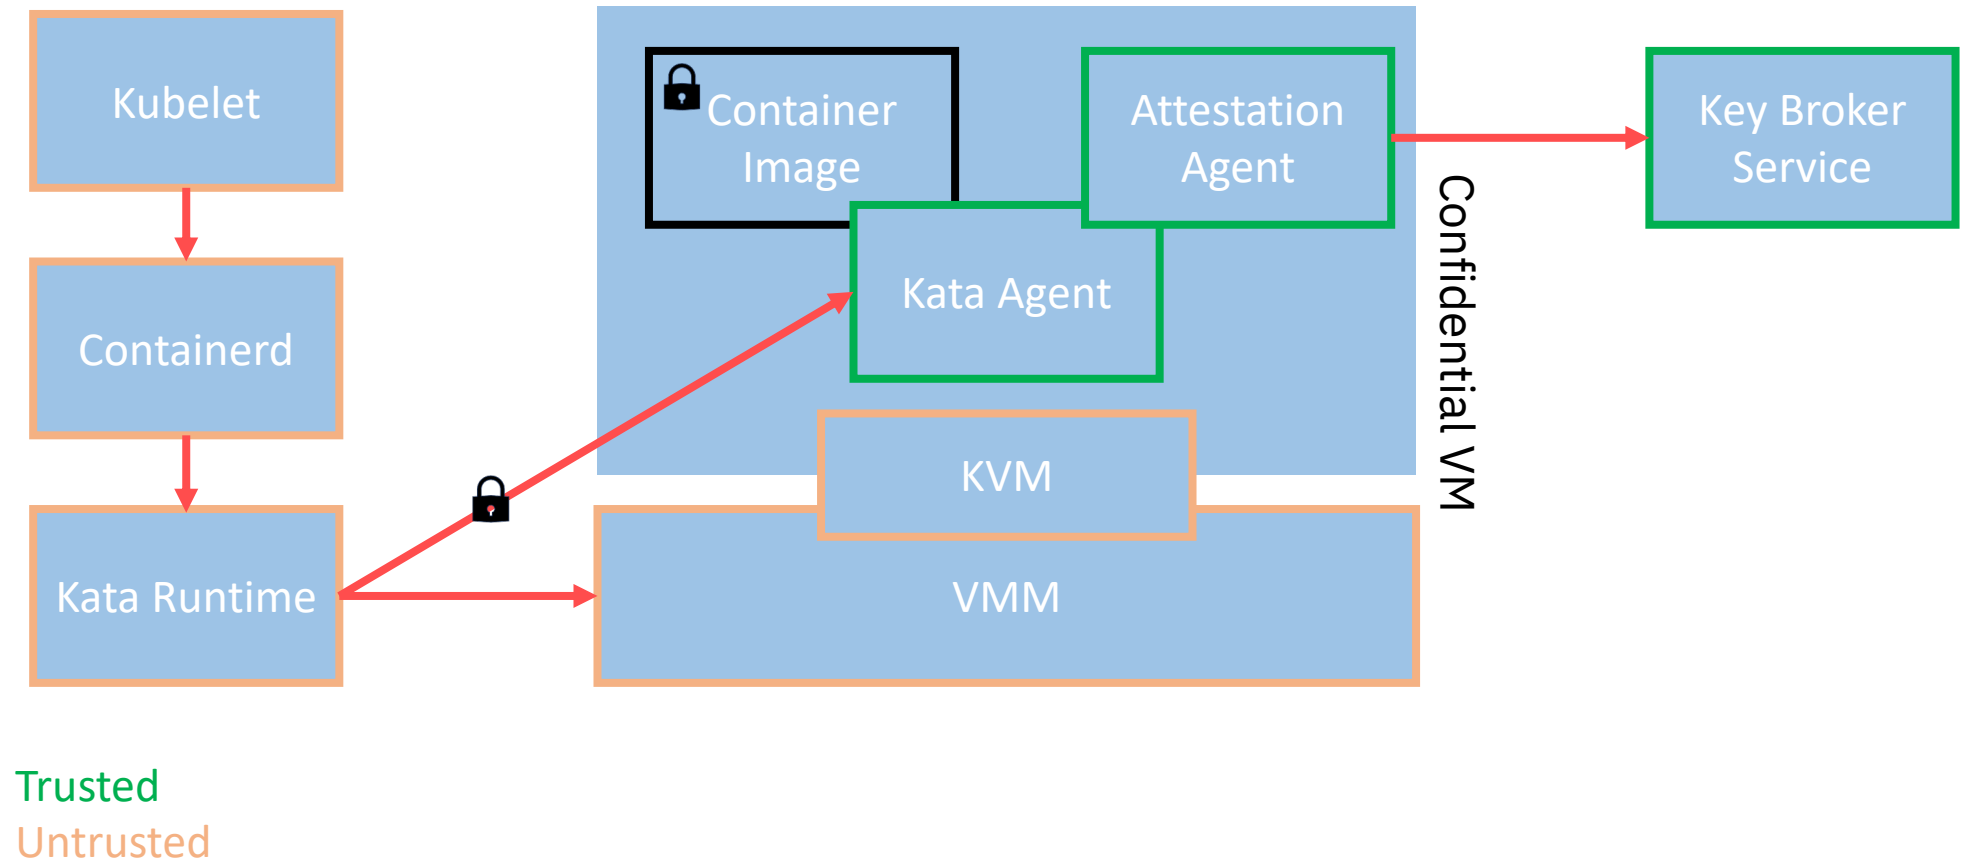

Container Breakout

- If a malicious container gains access to guest userspace
  - Can generate a valid attestation report
  - Request secrets from any KBS
- Generic guest image means that attestation reports are interchangeable between guests
- Containers execute arbitrary code by design
- Not actively exploitable
- Does the security of Confidential Containers reduce to the security of containers?

Attack in detail

1. Attacker creates a container that can execute userspace code
2. Attacker runs container with CoCo
3. Container gets nonce from target KBS
4. Container breaks out and gets attestation report with nonce and own public key
5. Container uses attestation report to get secret from target KBS

Solutions

- Revoke access to attestation reports
  - Phases of execution
  - The passport model
- Host data
- IMA
  - requires vTPM
- VMPLs

Host Data

- Field in attestation report set by the host prior to launch
- Put public key of KBS in host data
  - Binds the evidence to one KBS
  - Does not guarantee the identity of the KBS
  - Means that we can only attack one KBS at a time
- Turn on signature validation
  - Now the Attestation Agent must connect to a KBS to get the signature policy information
  - This KBS must match the public key in the host data
  - Now the only KBS we can connect to must be the one that signs the images

Notes

- This requires the target KBS to use signatures
  - Is this a reasonable assumption?
  - Image encryption is fundamentally optional and does not give the same guarantee
- We could bind the evidence to the workload
  - Hard to reuse evidence if host data specifies the workload
  - Breaks the decoupling of the workload and the guest
    - Aren’t signatures the best way to measure the workload?

What about SEV(-ES)

- Connection to KBS is made from the host
  - Very hard to regulate
- Connection cannot be revoked inside guest
- VM can only connect to one KBS at a time

Conclusions

- The capability of generating valid evidence needs to be protected
- Consider what will happen if the protections fail

- Don’t rely too much on one trust model