Protecting Sensitive Data in Container Checkpoints

Radostin Stoyanov - PhD Student, Scientific Computing Group

Collaboration with Adrian Reber, Senior Principal Software Engineer

Supervisor: Prof. Wes Armour

Container Checkpointing

kubernetes

podman

docker

cri-o

containerd

RUNC

crun

youki

LXC/LXD

CRIU

Ubuntu

fedora

debian

CentOS

alpine Linux

archlinux

Checkpoint Encryption

End-to-end encryption of sensitive checkpoint data

Checkpoint Encryption

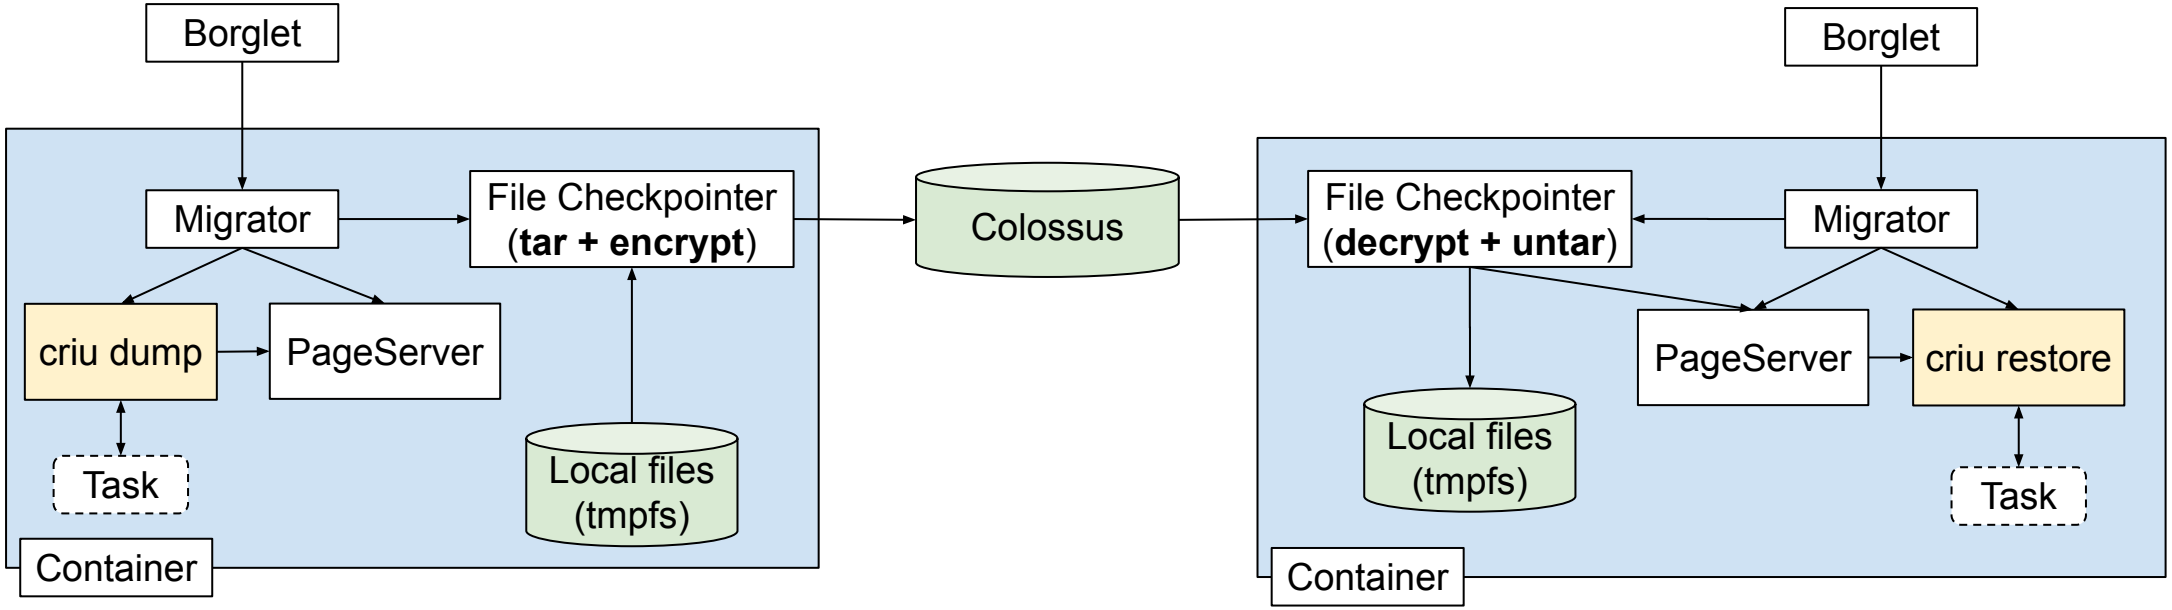

Andy Tucker (Google). *Task Migration at Google Using CRIU*. Linux Plumbers Conference. 2018

Checkpoint Encryption

- Easy to use across different distributions & container engines
- Easy to integrate with Kubernetes
- Support for iterative checkpointing (`pre-dump` & `--auto-dedup`)
- Public-key infrastructure & key management
- Low performance and space overhead

Checkpoint Encryption

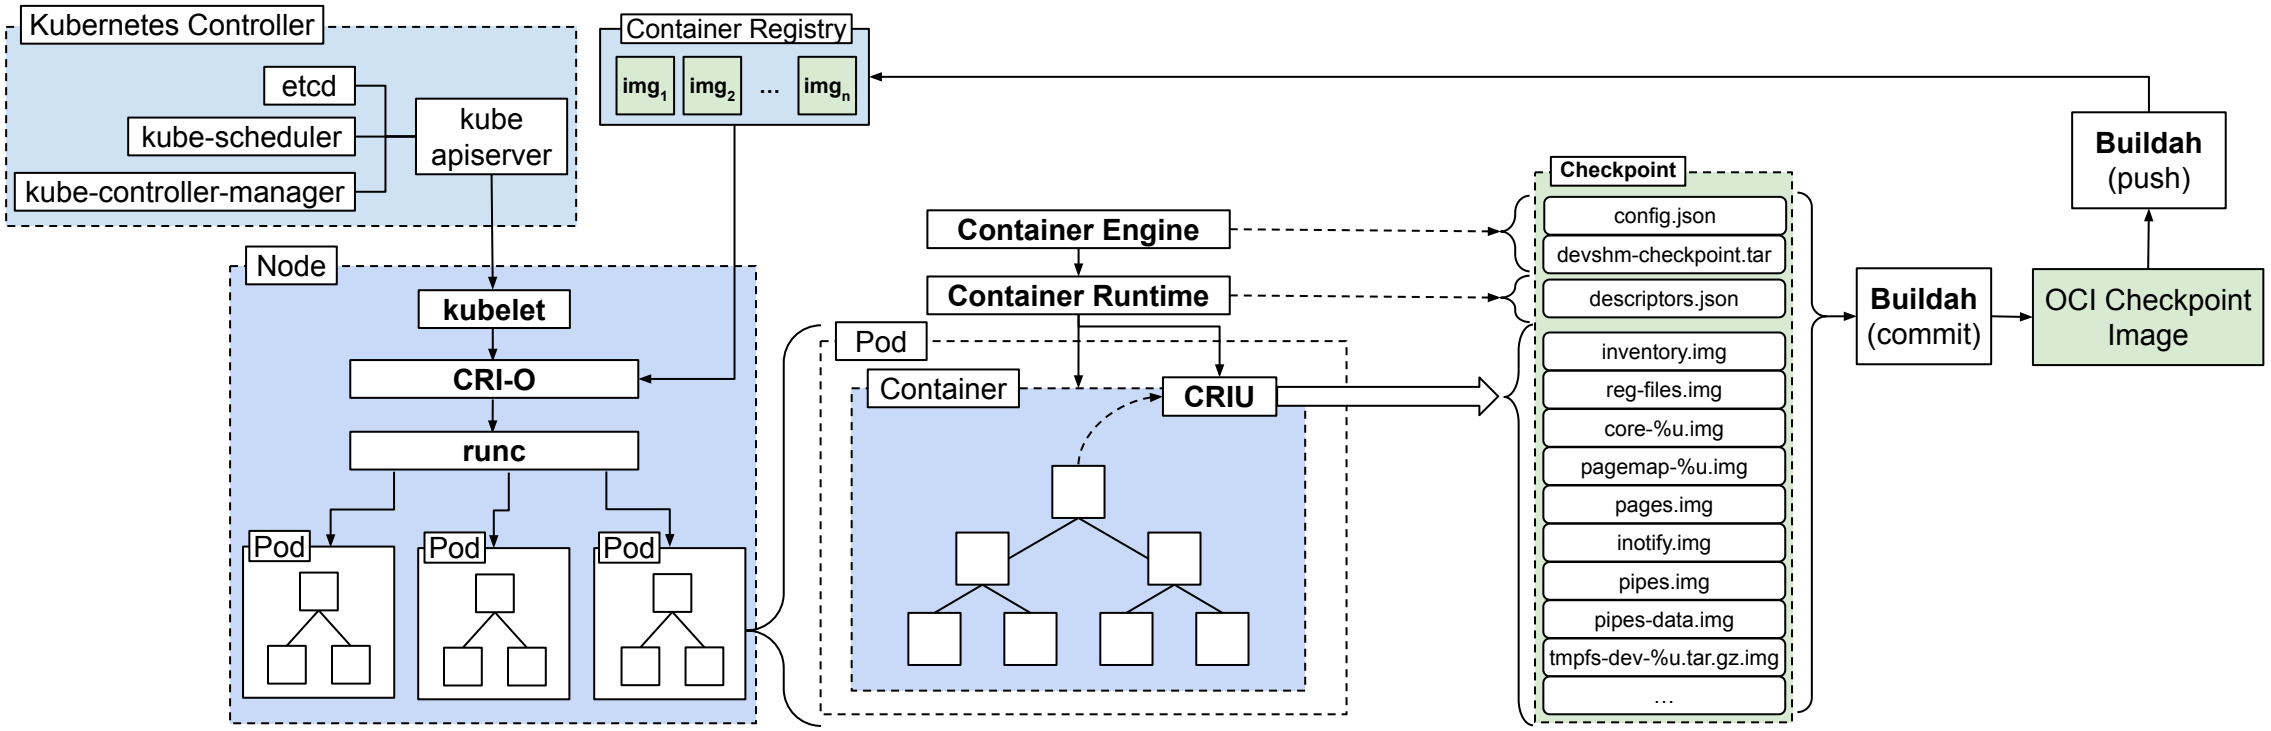

Checkpoint Encryption

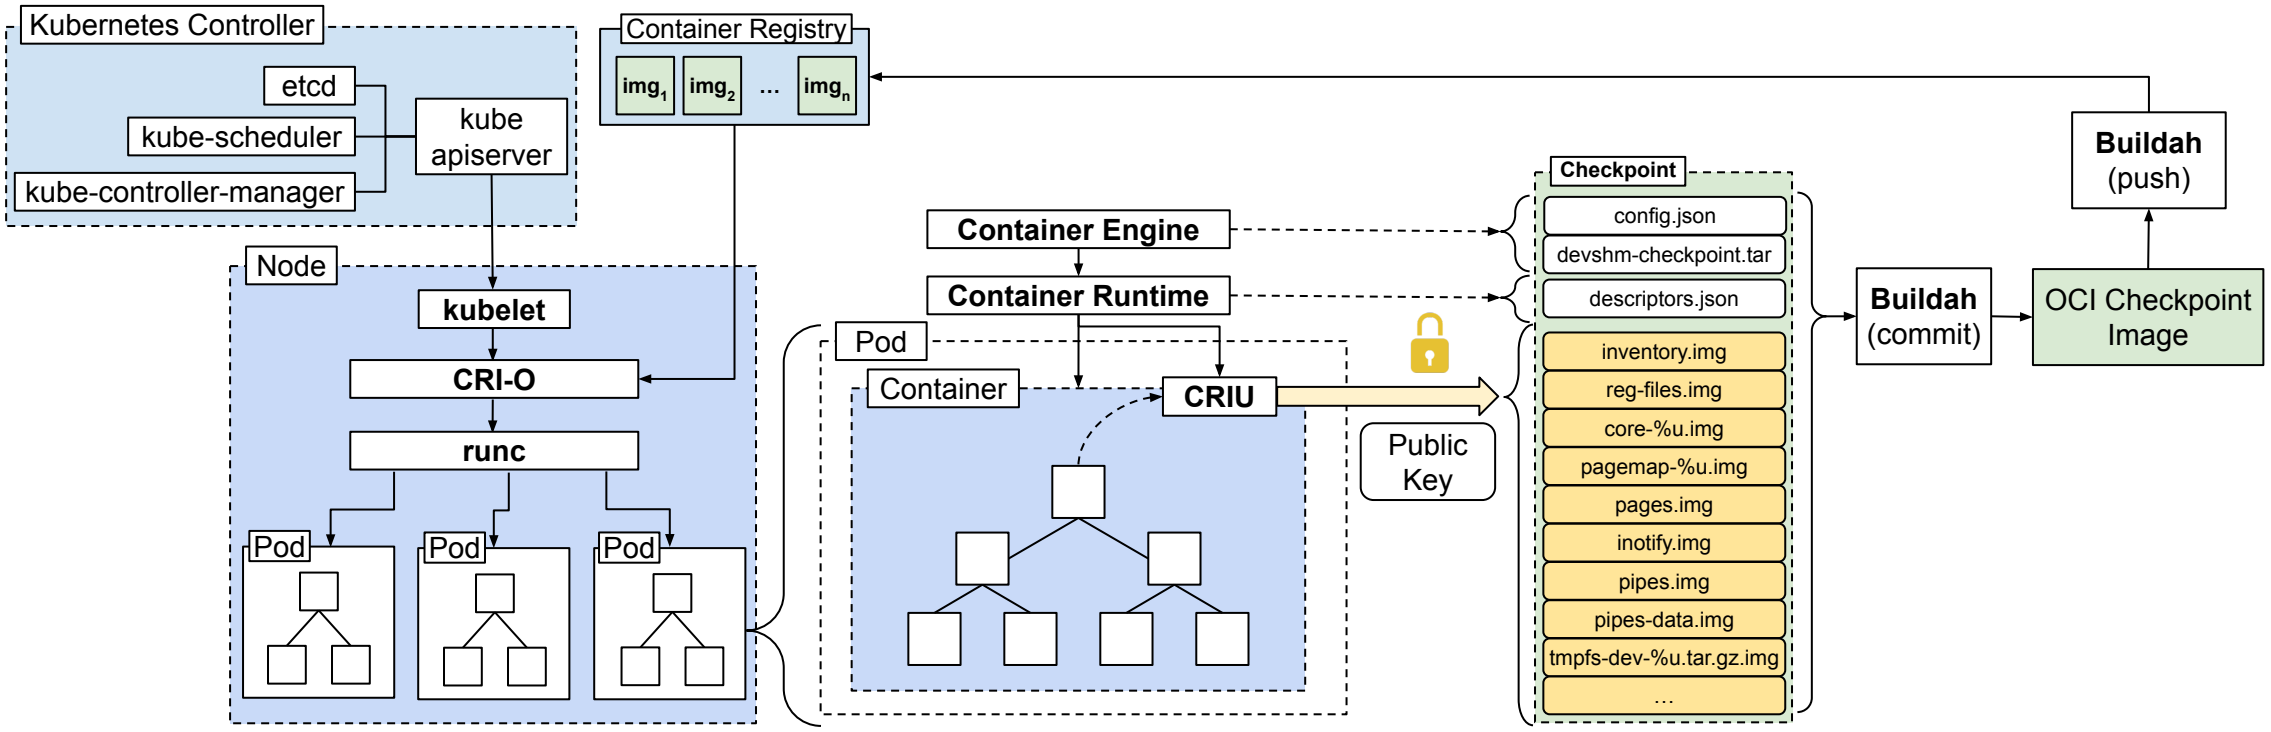

Adding support for built-in encryption


```
criu dump --encrypt ...
```

- GnuTLS provides the necessary crypto primitives
  - Seamless integration with container engines
  - Code reuse with different images
- Integration with CRIT (decode encrypted images)
- Optimized for performance and space efficiency
- End-to-end encryption – sensitive data is always encrypted

Protobuf Images

Encrypting CRIU-specific images in protocol buffer format

ChaCha20-Poly1305 AEAD

- Appropriate for small messages
- Authenticated Encryption with Associated Data
  - Provides confidentiality, integrity, and authenticity
  - 256-bit key, 96-bit nonce, 128-bit tag
- Great efficiency in software implementations compared to AES-GCM

Protobuf Images

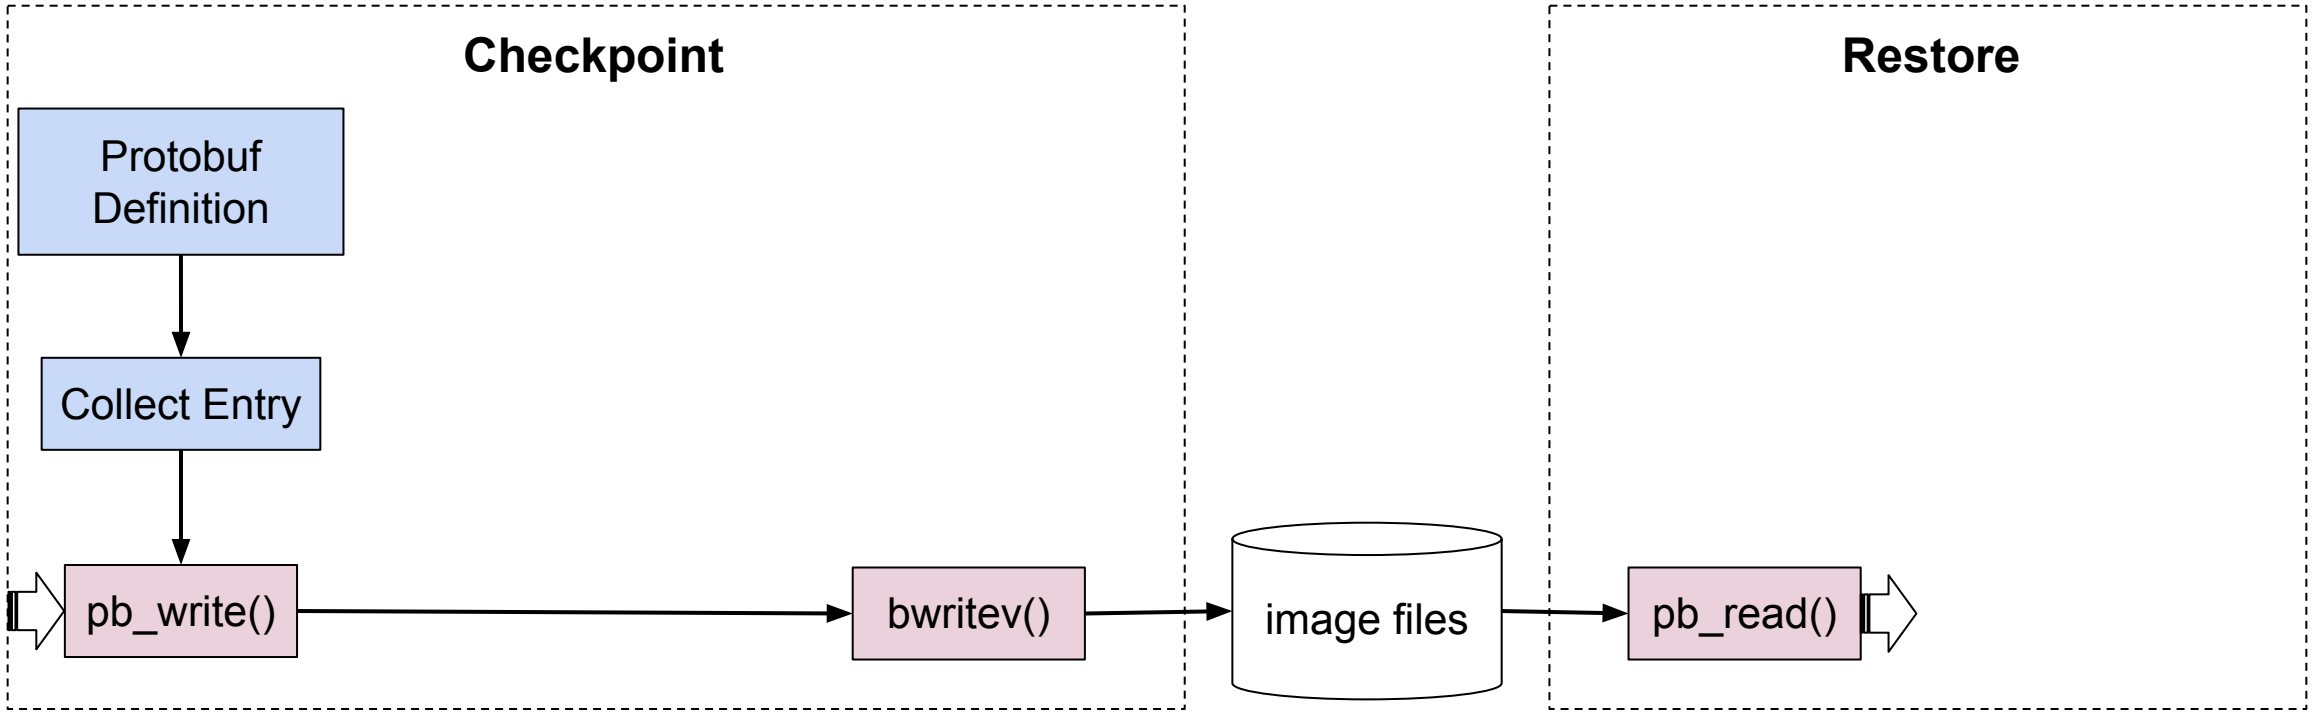

Protobuf Images

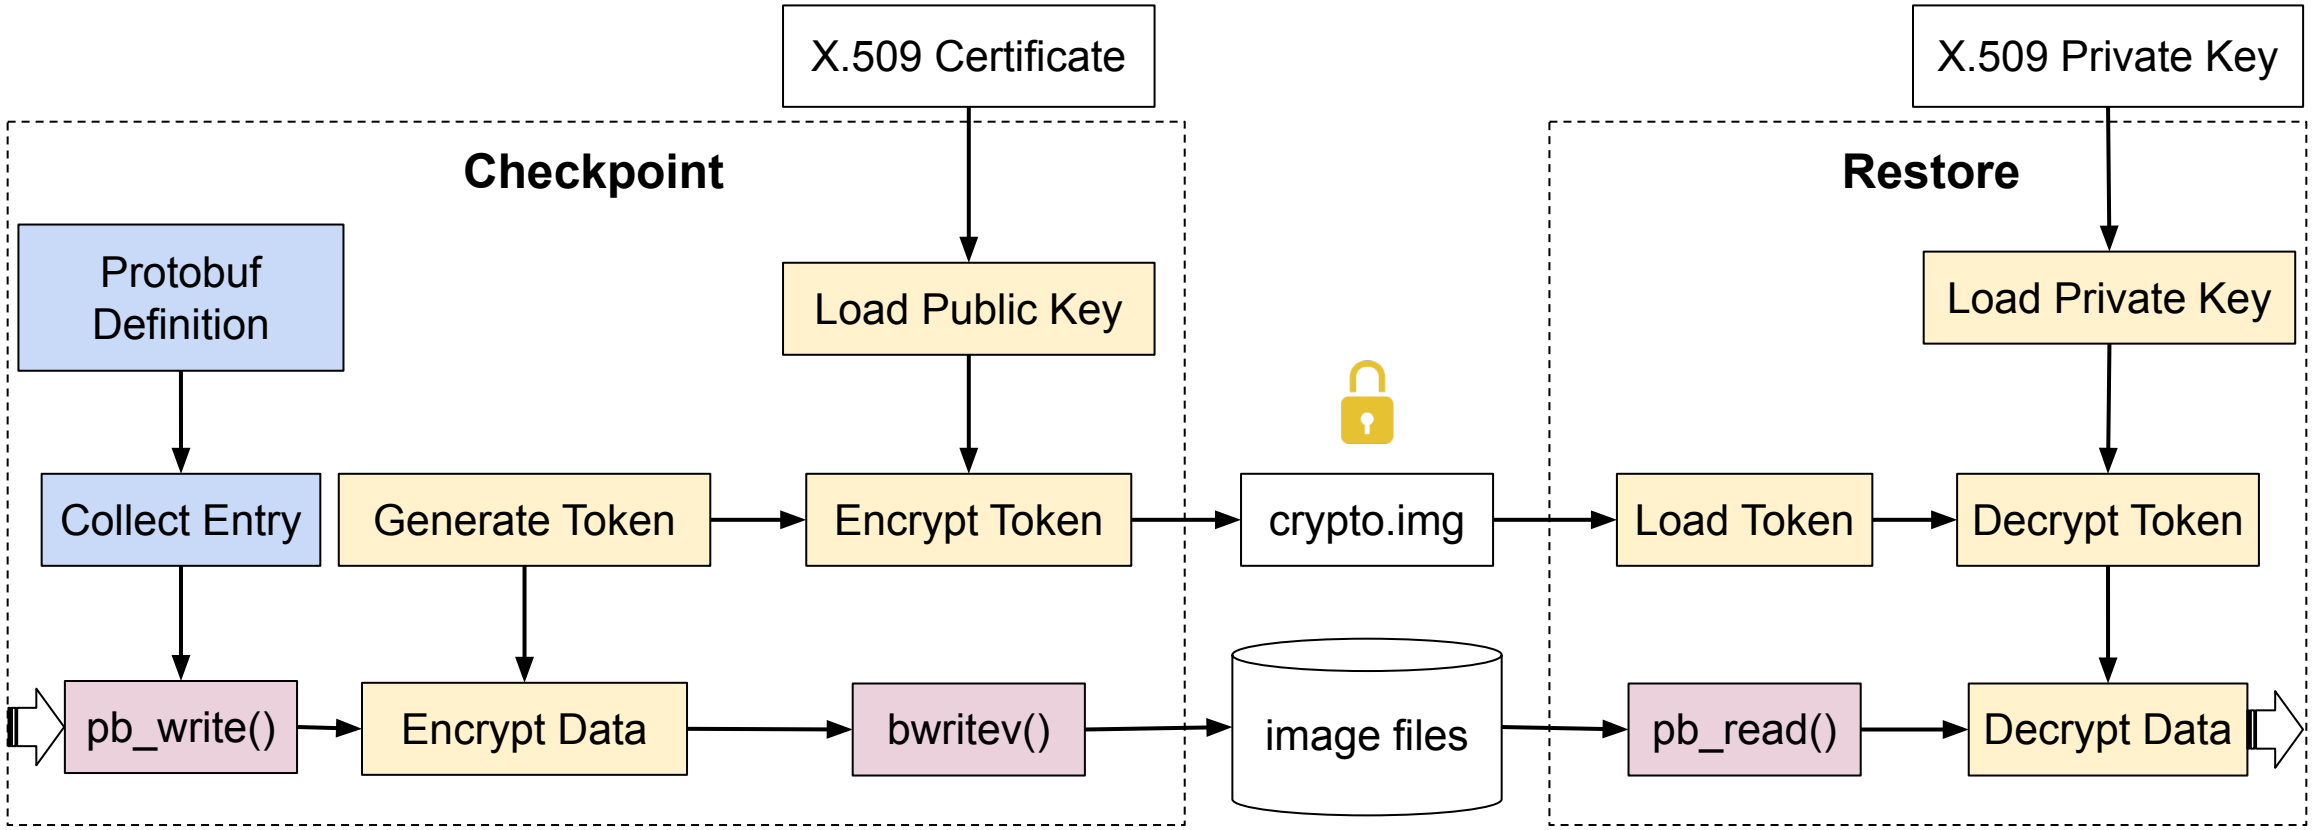

Raw Images

Encrypting data collected with the help of external tools

Raw Images

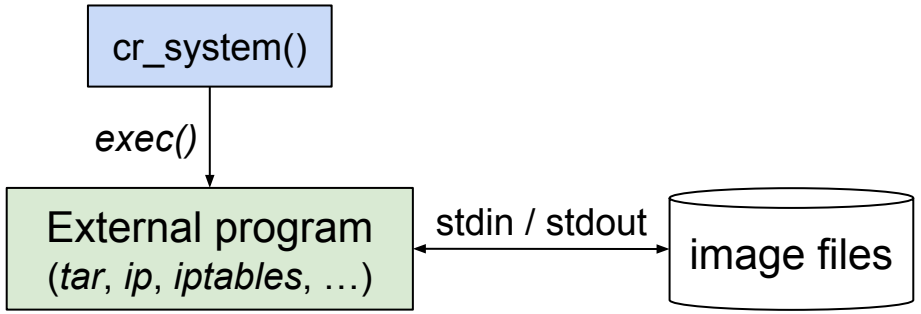

https://criu.org/Images#Raw_images

Raw Images

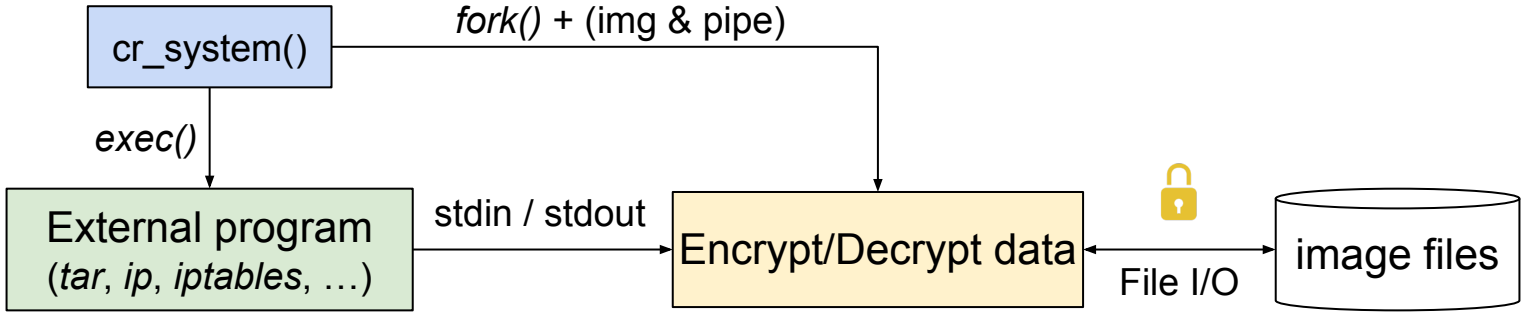

https://criu.org/Images#Raw_images

Memory Pages

Encrypting the contents of individual pages (4k)

AES-XTS

- XEX-based tweaked-codebook mode with ciphertext stealing (XTS)
  - 256-bit key + 256-bit tweak key
  - Initialization vector (reduce space overhead)
- Preferred for block/disk encryption
- Hardware acceleration
  - ~5x faster than software
  - requires gnutls v3.6.14 or newer

Suggested by Daiki Ueno (GnuTLS team)

Memory Pages

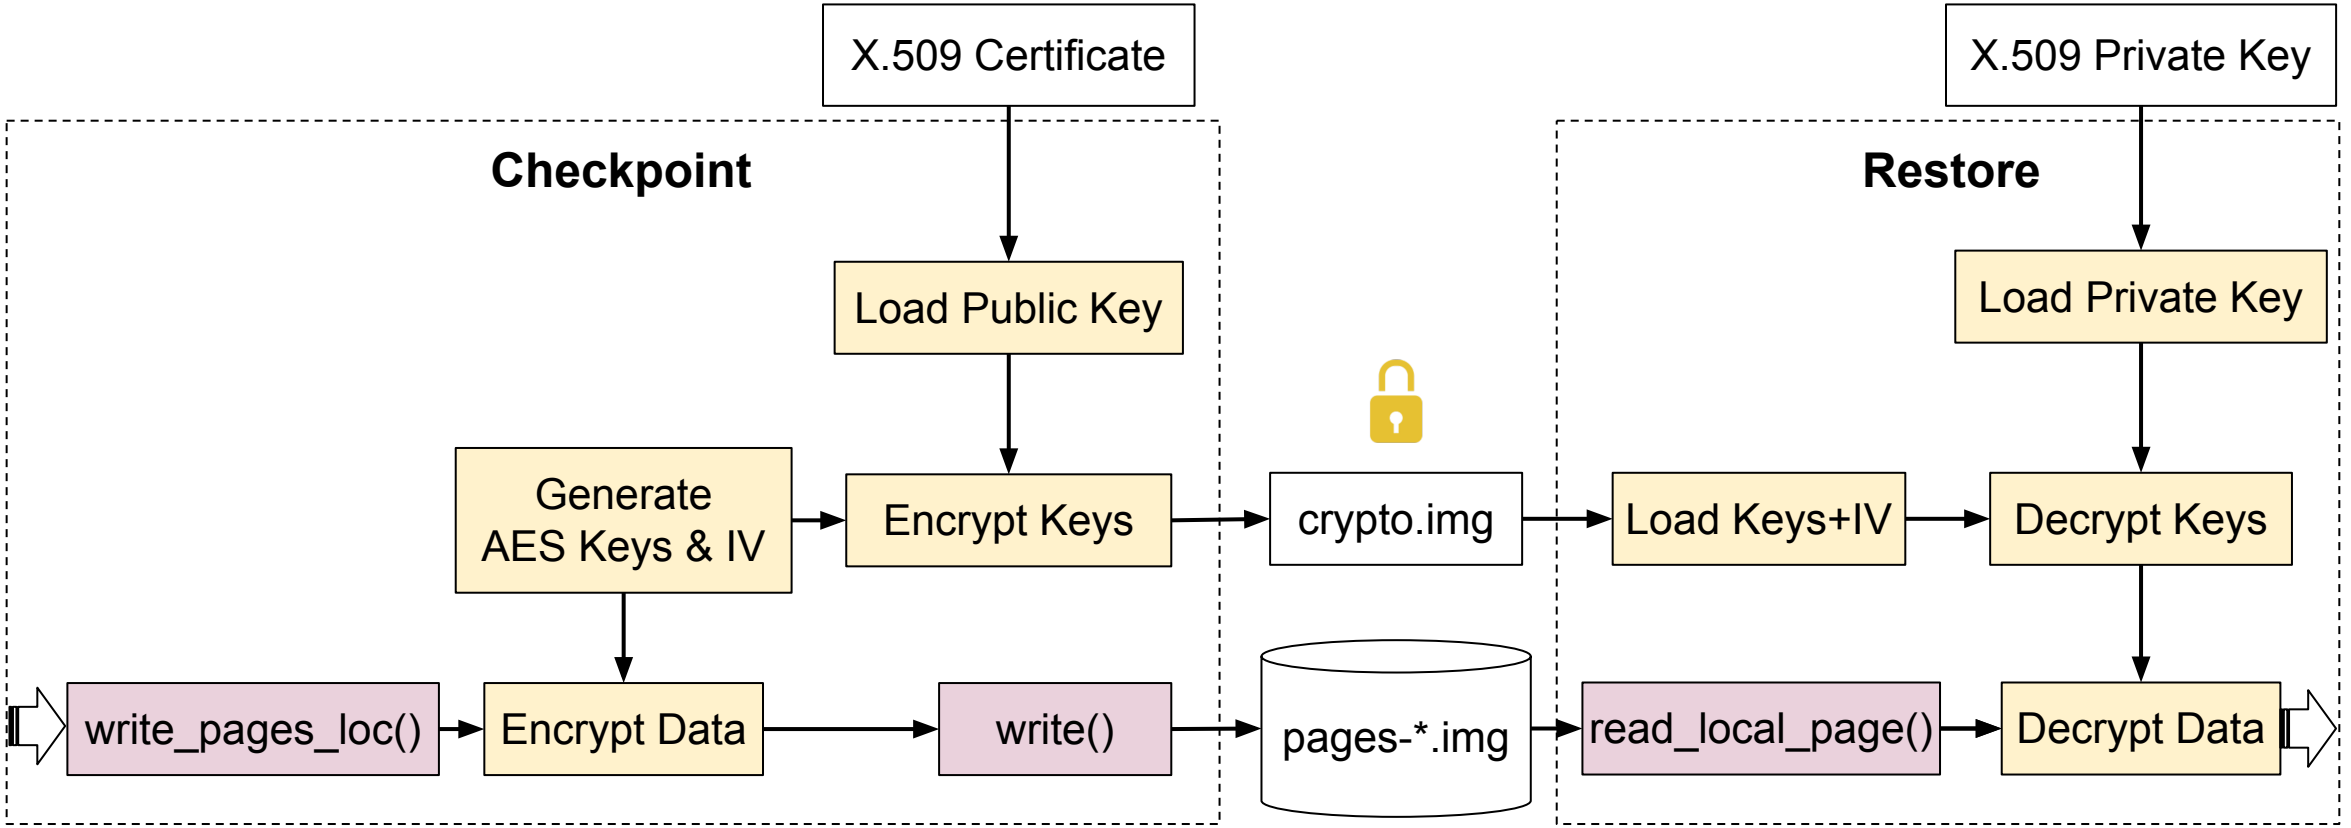

Memory Pages

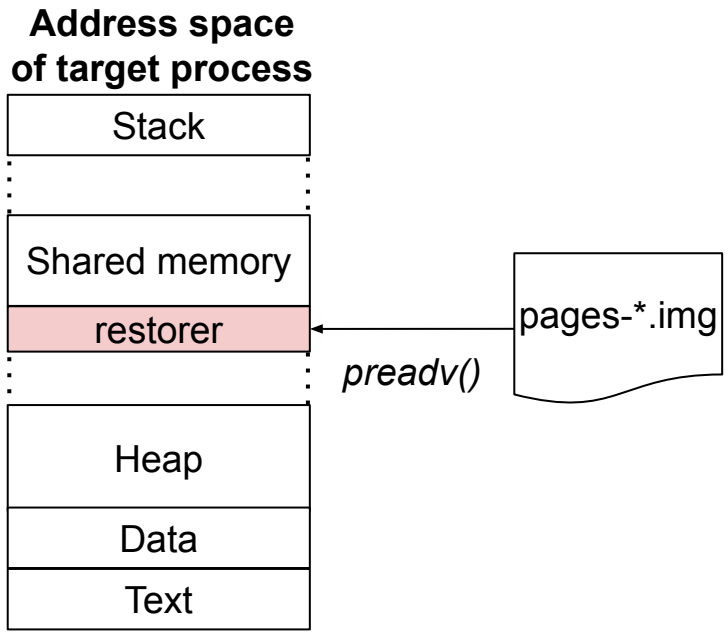

https://criu.org/copy-on-write_memory

Memory Pages

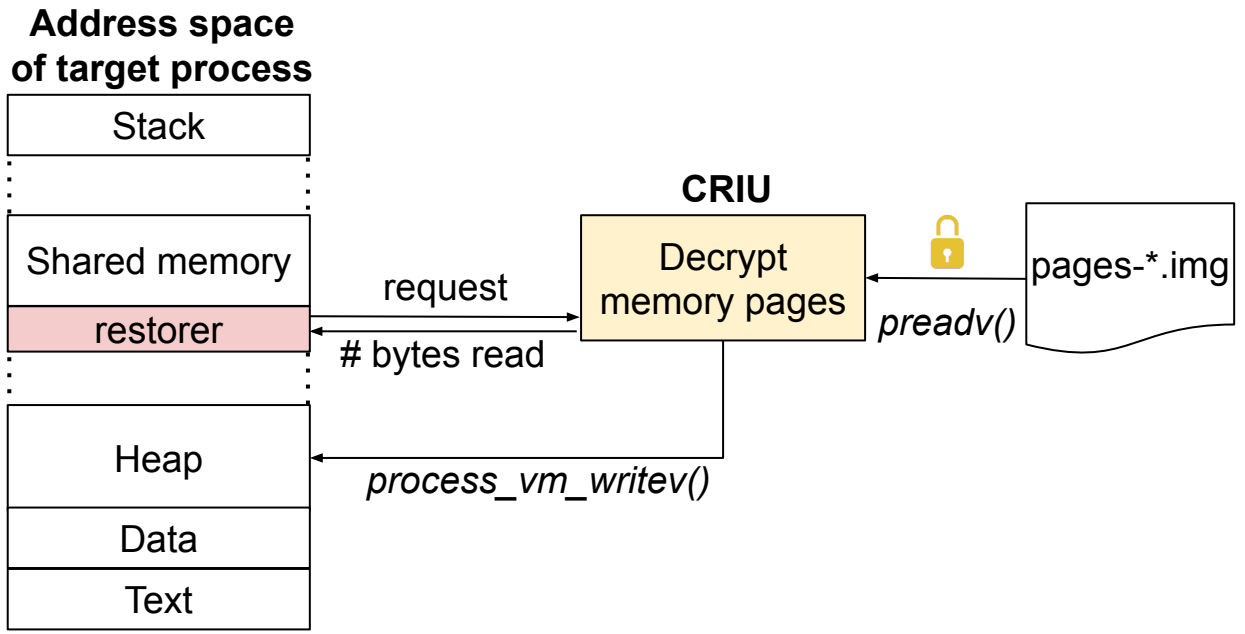

CRIT

Decoding encrypted images

Using CRIT with encrypted images

- *cryptography* Python module provides all necessary cryptographic primitives
- If *cipher.img* exists, load *token* value
- If *token* has been loaded, use it to decrypt image data
- `--tls-key` can be used to specify path to private key (PEM file)

ZDTM

Running tests with encrypted images

ZDTM tests with encryption

- Command-line option to run tests with encryption
  - Example:

```
zdtm.py run -t zdtm/static/busyloop00 --encrypt
```

- PKI certificate and private key
  - `./test/pki/*.pem`

Demo

Performance Evaluation

DEPARTMENT OF ENGINEERING SCIENCE

UNIVERSITY OF OXFORD

stress-ng

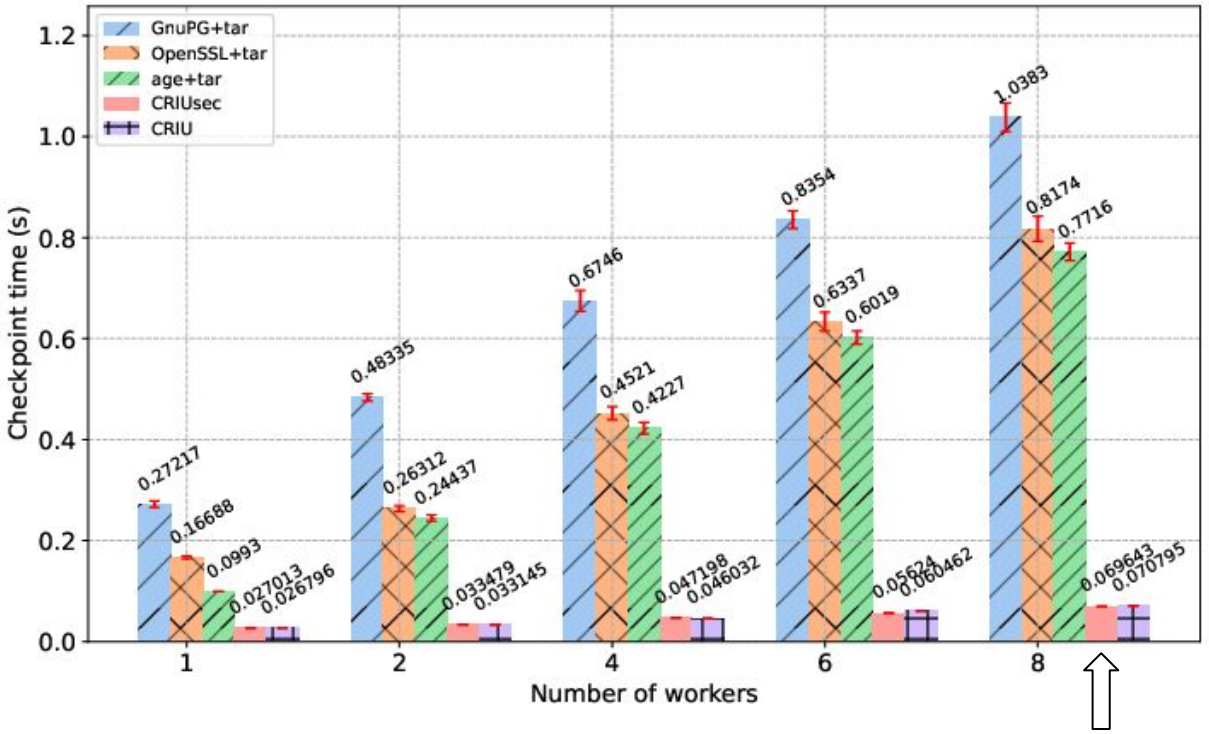

memhog

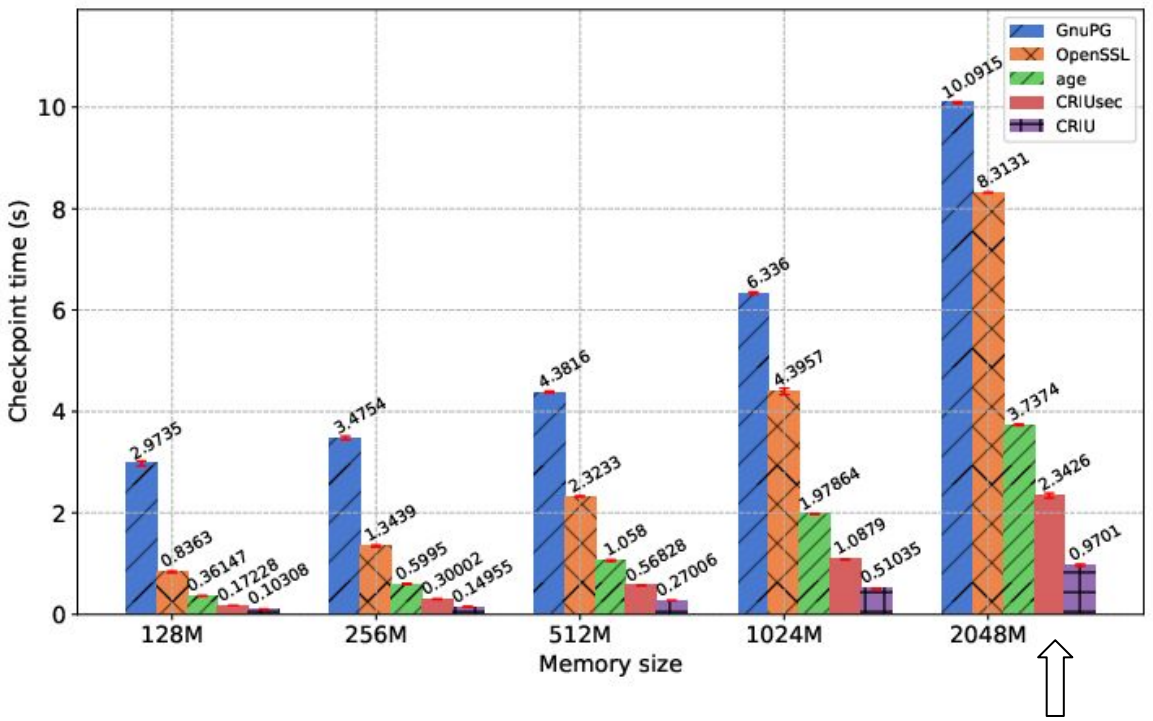

Conclusion

- Encryption support for protobuf, raw images and memory pages
- ZDTM tests with encrypted images
- CRIT support for decoding encrypted images

- Future work
  - checkpointctl support for encrypted images
  - Iterative checkpointing & memory deduplication
  - Adaptive compression

Questions?

[RFC] Add support for encrypted images

https://github.com/checkpoint-restore/criu/pull/2297