Ginseng, the Learning TEE

Fast, Confidential Machine Learning in FPGA Enclaves

Nick Hynes | Oasis Labs

Ideal: data providers pool data to train a large, complex model

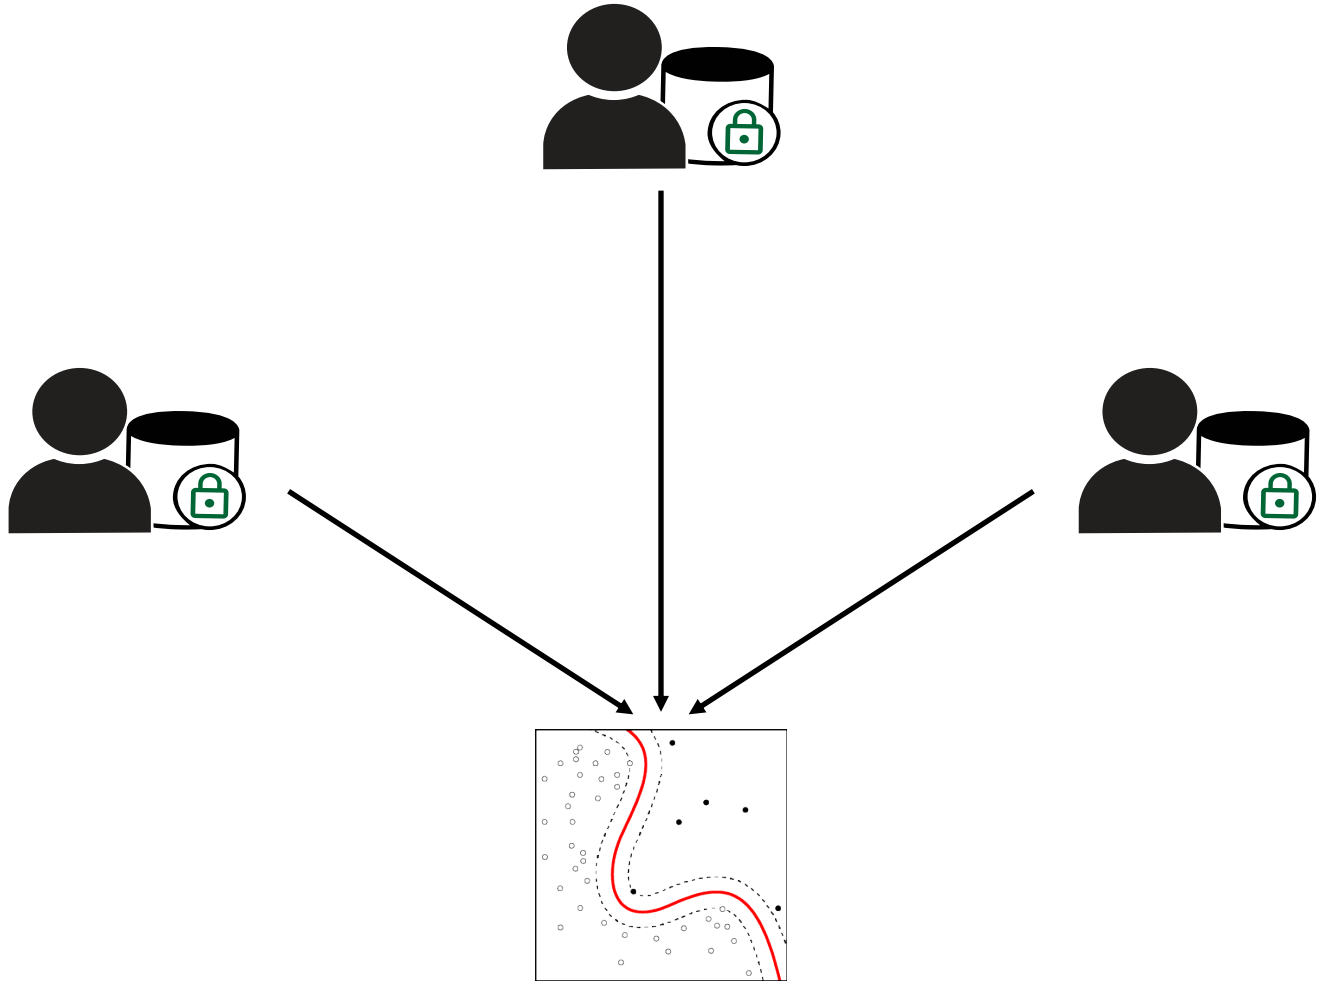

Ideal: data providers pool data to train a large, complex model

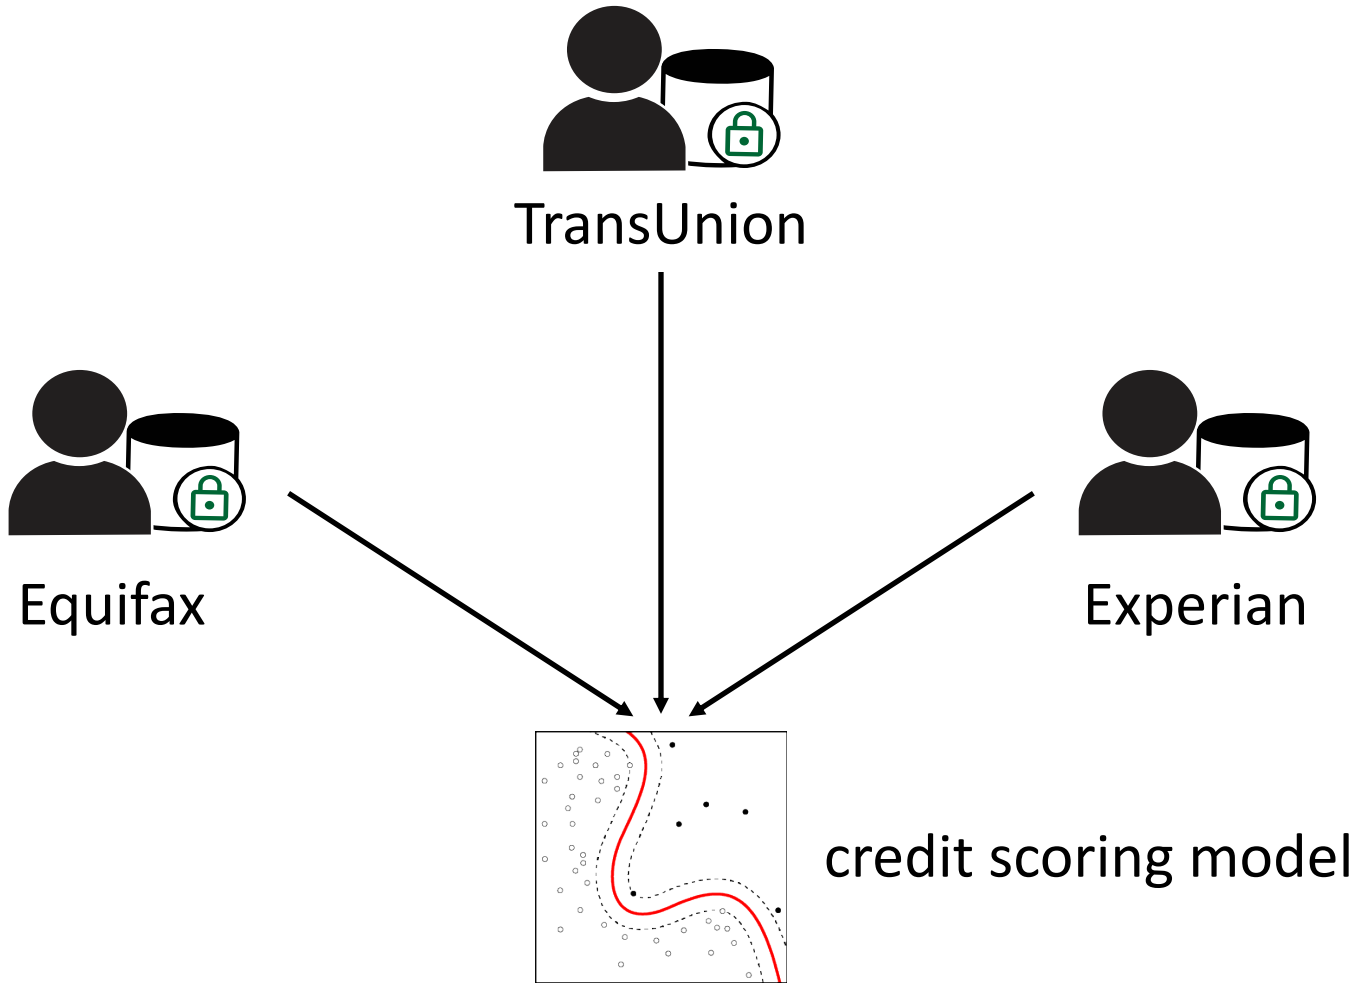

Ideal: data providers pool data to train a large, complex model

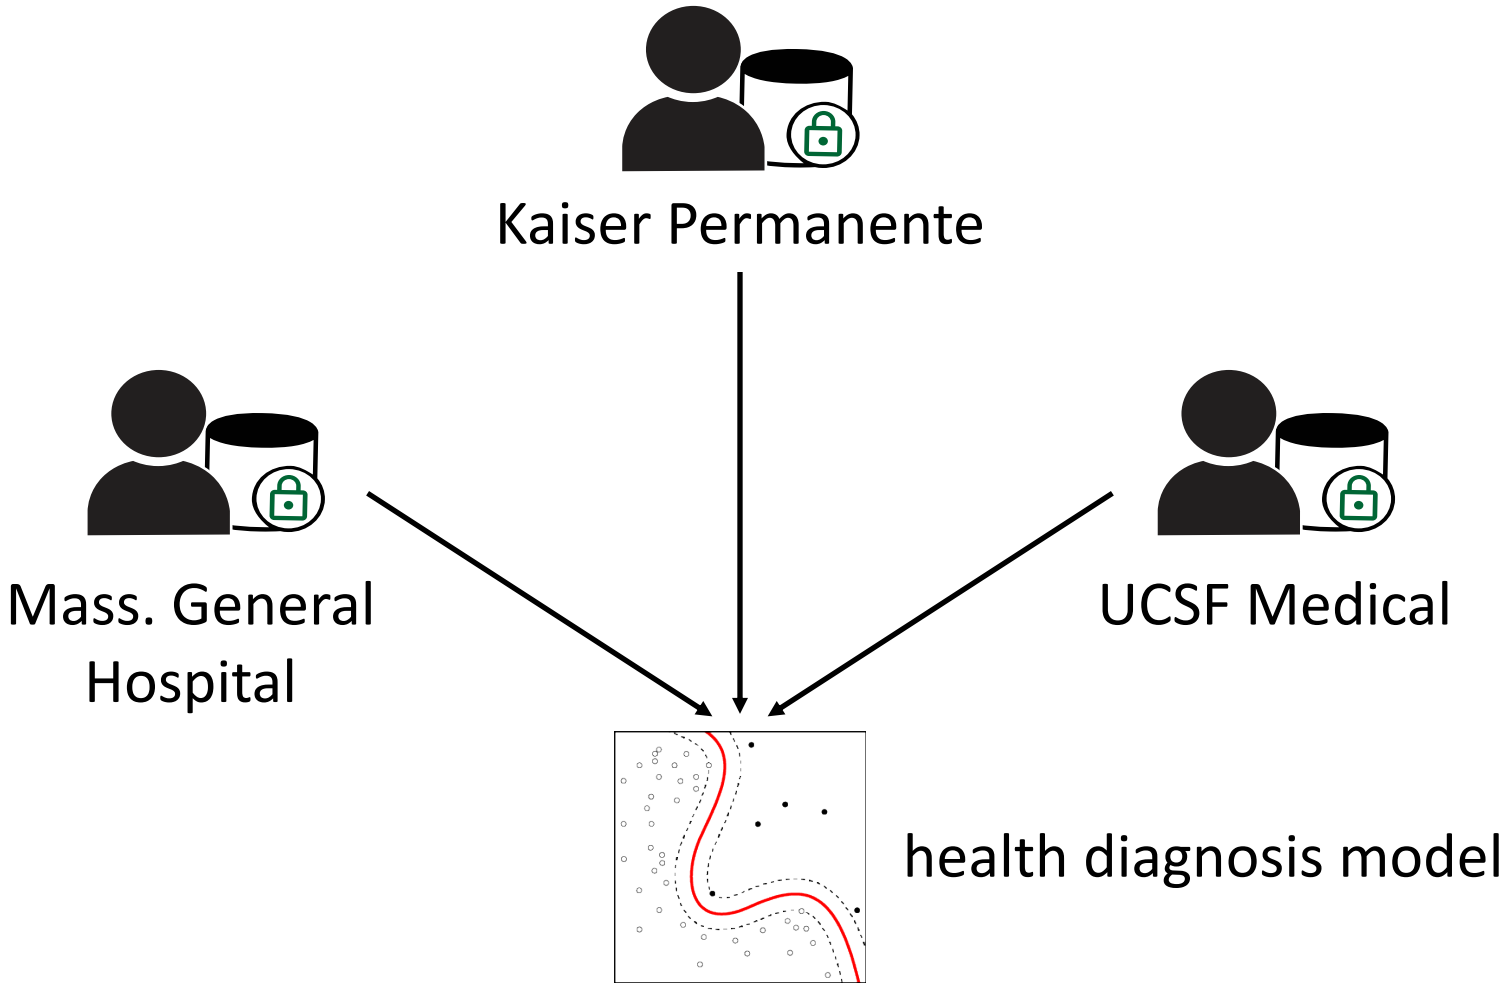

Ideal: data providers pool data to train a large, complex model

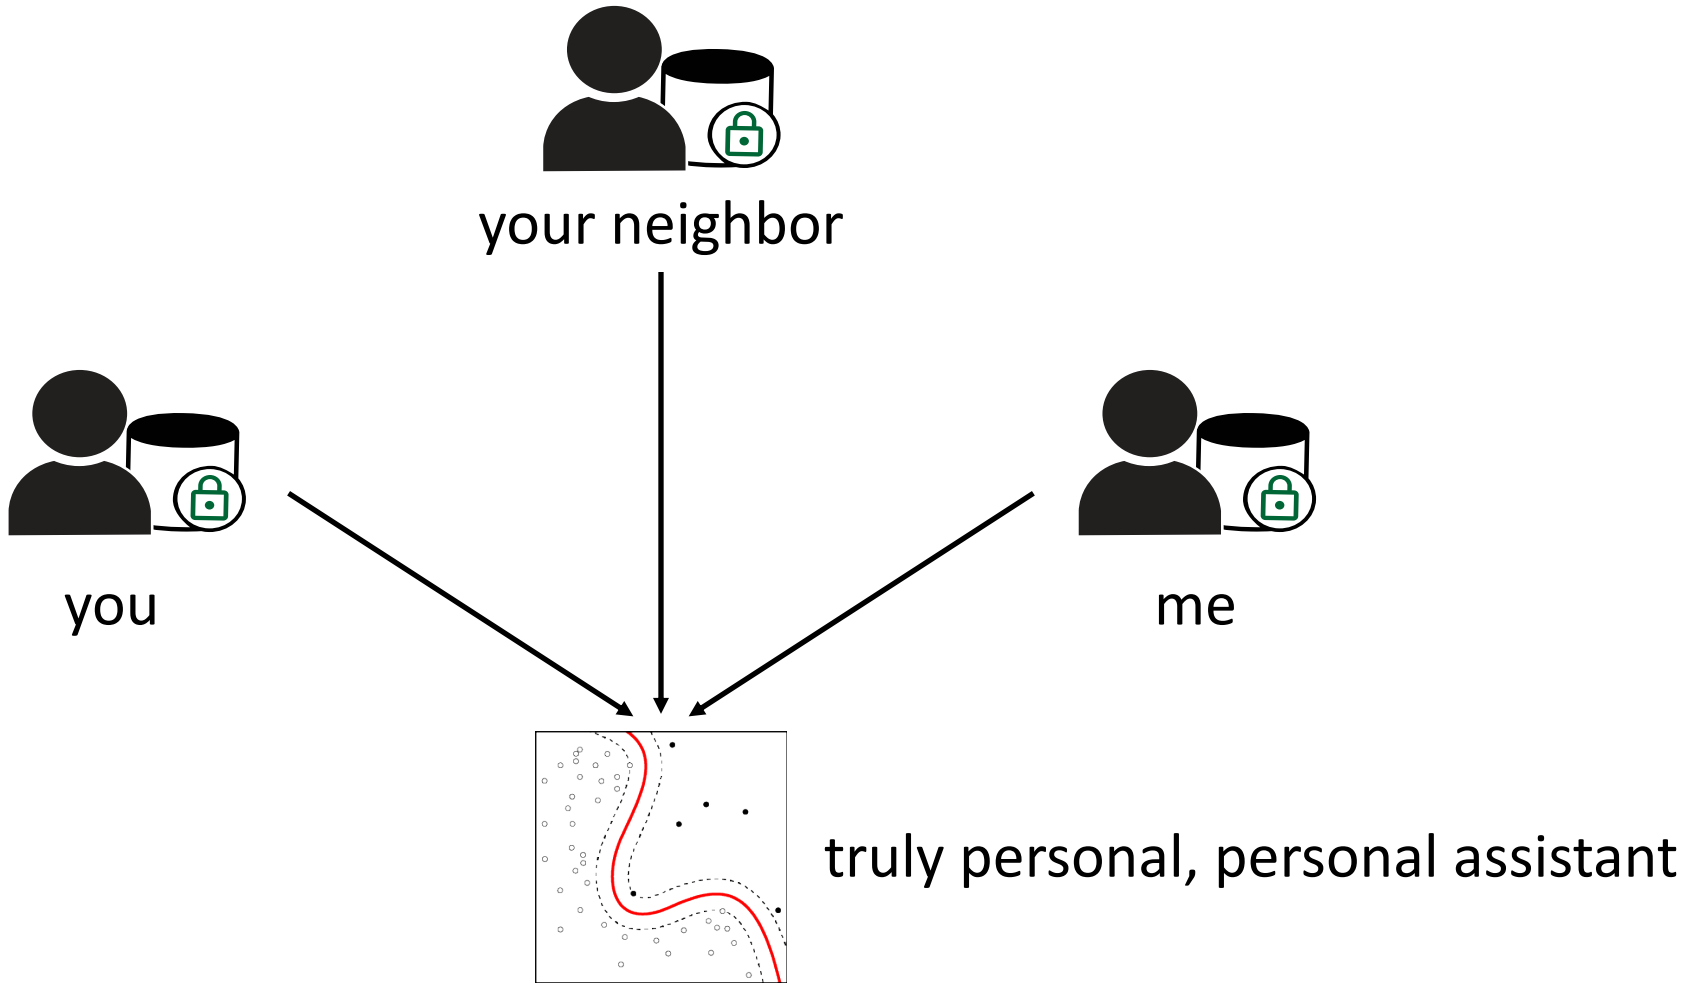

Reality: data providers are mutually distrusting!

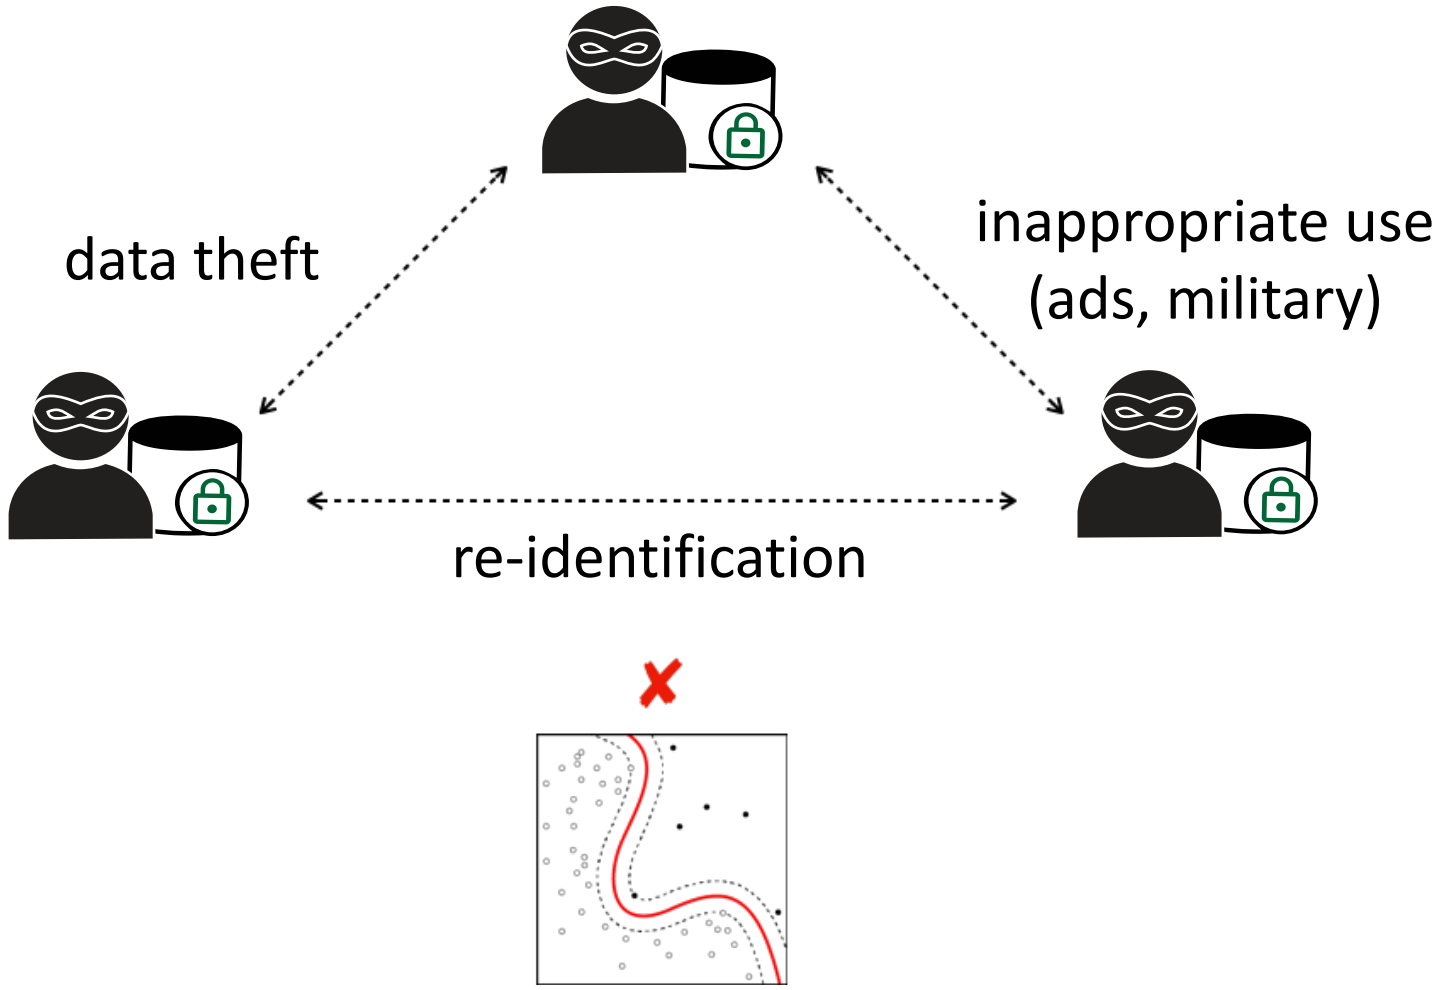

Solution: cooperation via a trusted third party (i.e. enclave)

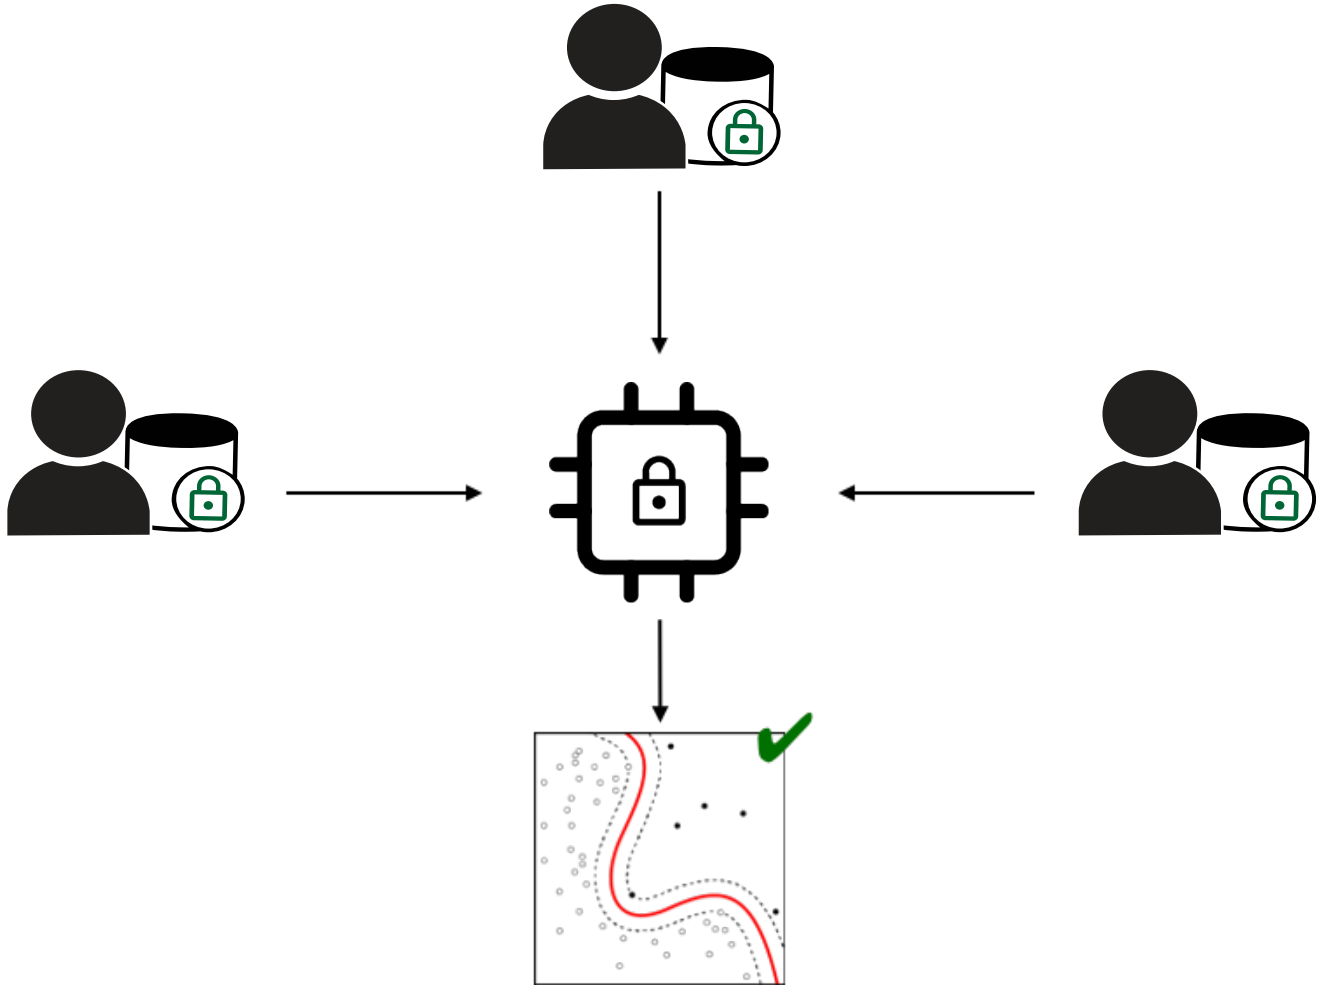

What about CPU Enclaves?

Performance of VGG-9 on CIFAR (32x32 RGB images)

	img/s (training)	img/s (inference)
Myelin [1]	21 img/s	496 img/s
Chiron (4 enclaves) [2]	25 img/s	–
non-private CPU	42 img/s	1119 img/s

[1] *Efficient Deep Learning on Multi-Source Private Data*. N. Hynes, R. Cheng, D. Song. Arxiv 2018
[2] *Chiron: Privacy-preserving machine learning as a service*. T. Hunt, C. Song, R. Shokri, V. Shmatikov, and E. Witchel. Arxiv 2018
[3] *Graviton: Trusted Execution Environments on GPUs*. S. Volos, K. Vaswani. OSDI 2018

What about CPU Enclaves?

Performance of VGG-9 on CIFAR (32x32 RGB images)

	img/s (training)	img/s (inference)
Myelin [1]	21 img/s	496 img/s
Chiron (4 enclaves) [2]	25 img/s	–
non-private CPU	42 img/s	1119 img/s
private GPU: Graviton [3]	>1500 img/s	>10,000 img/s

[1] *Efficient Deep Learning on Multi-Source Private Data*. N. Hynes, R. Cheng, D. Song. Arxiv 2018
[2] *Chiron: Privacy-preserving machine learning as a service*. T. Hunt, C. Song, R. Shokri, V. Shmatikov, and E. Witchel. Arxiv 2018
[3] *Graviton: Trusted Execution Environments on GPUs*. S. Volos, K. Vaswani. OSDI 2018

Ginseng, the Learning TEE

FPGA-based ML accelerator

1. Start with a tensor accelerator framework (e.g., VTA [4])
2. Bolt on a Tensor Encryption Core (TEC)
3. Add remote attestation hardware (PUF, RNG)
4. Distribute with a lightweight, secure unikernel

End result: a speedy end-to-end private ML pipeline

[4] *A Hardware-Software Blueprint for Flexible Deep Learning Specialization*. T. Moreau, et al. Arxiv 2019

Ginseng, the Learning TEE

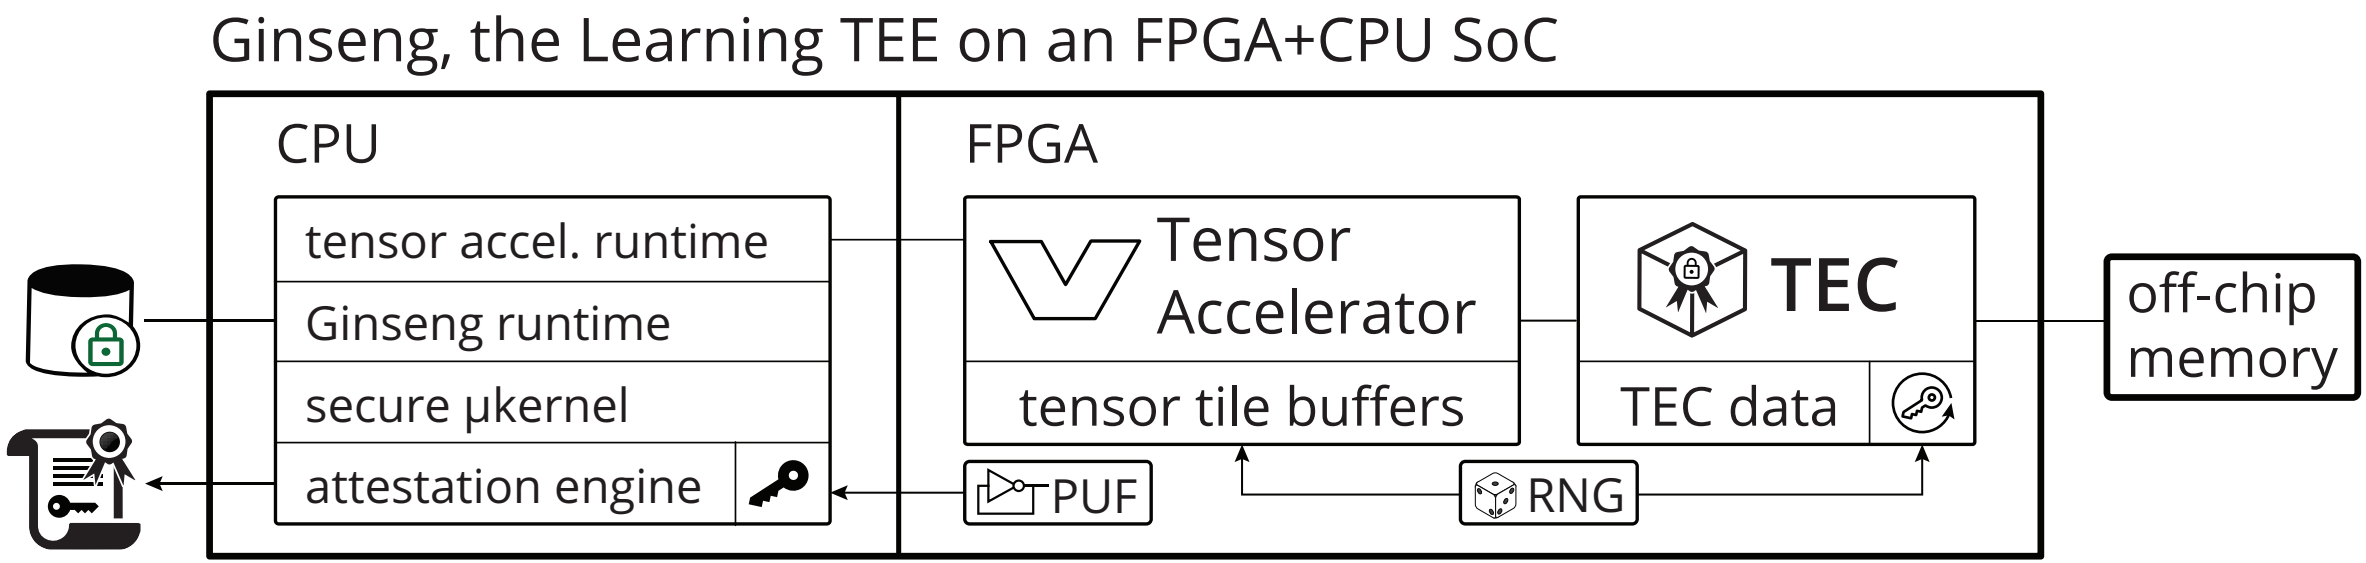

Ginseng, the Learning TEE

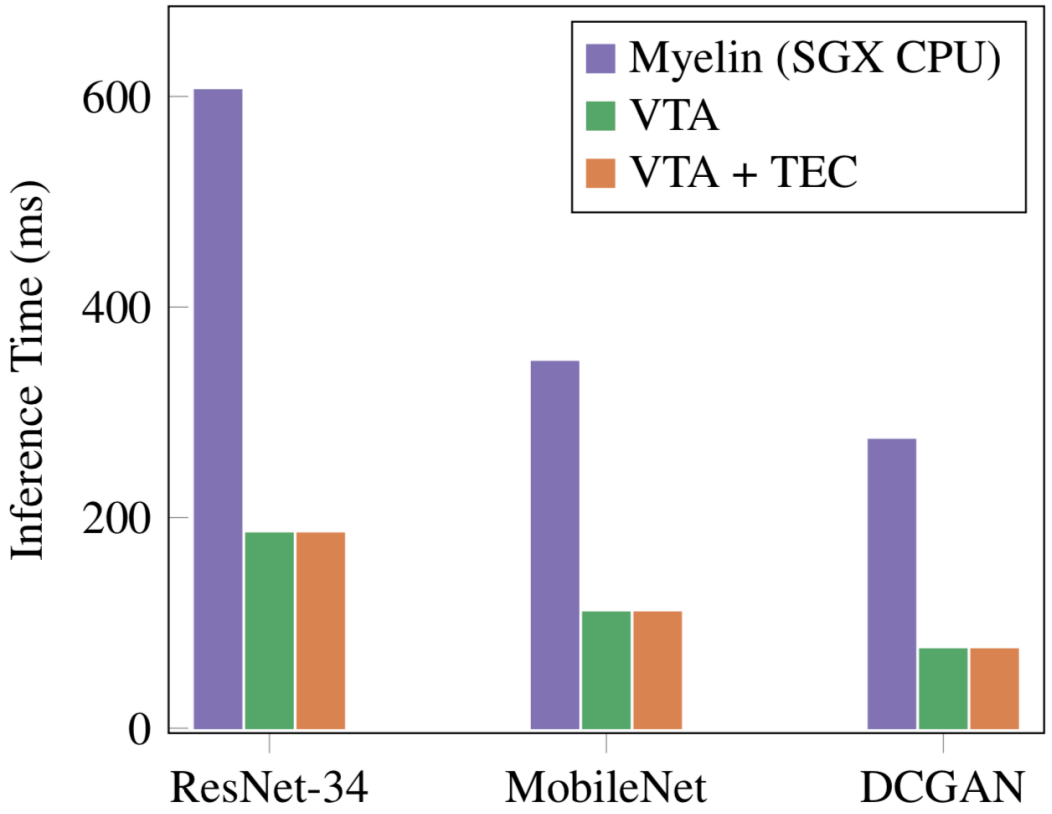

Sterling: A Privacy-Preserving Data Marketplace

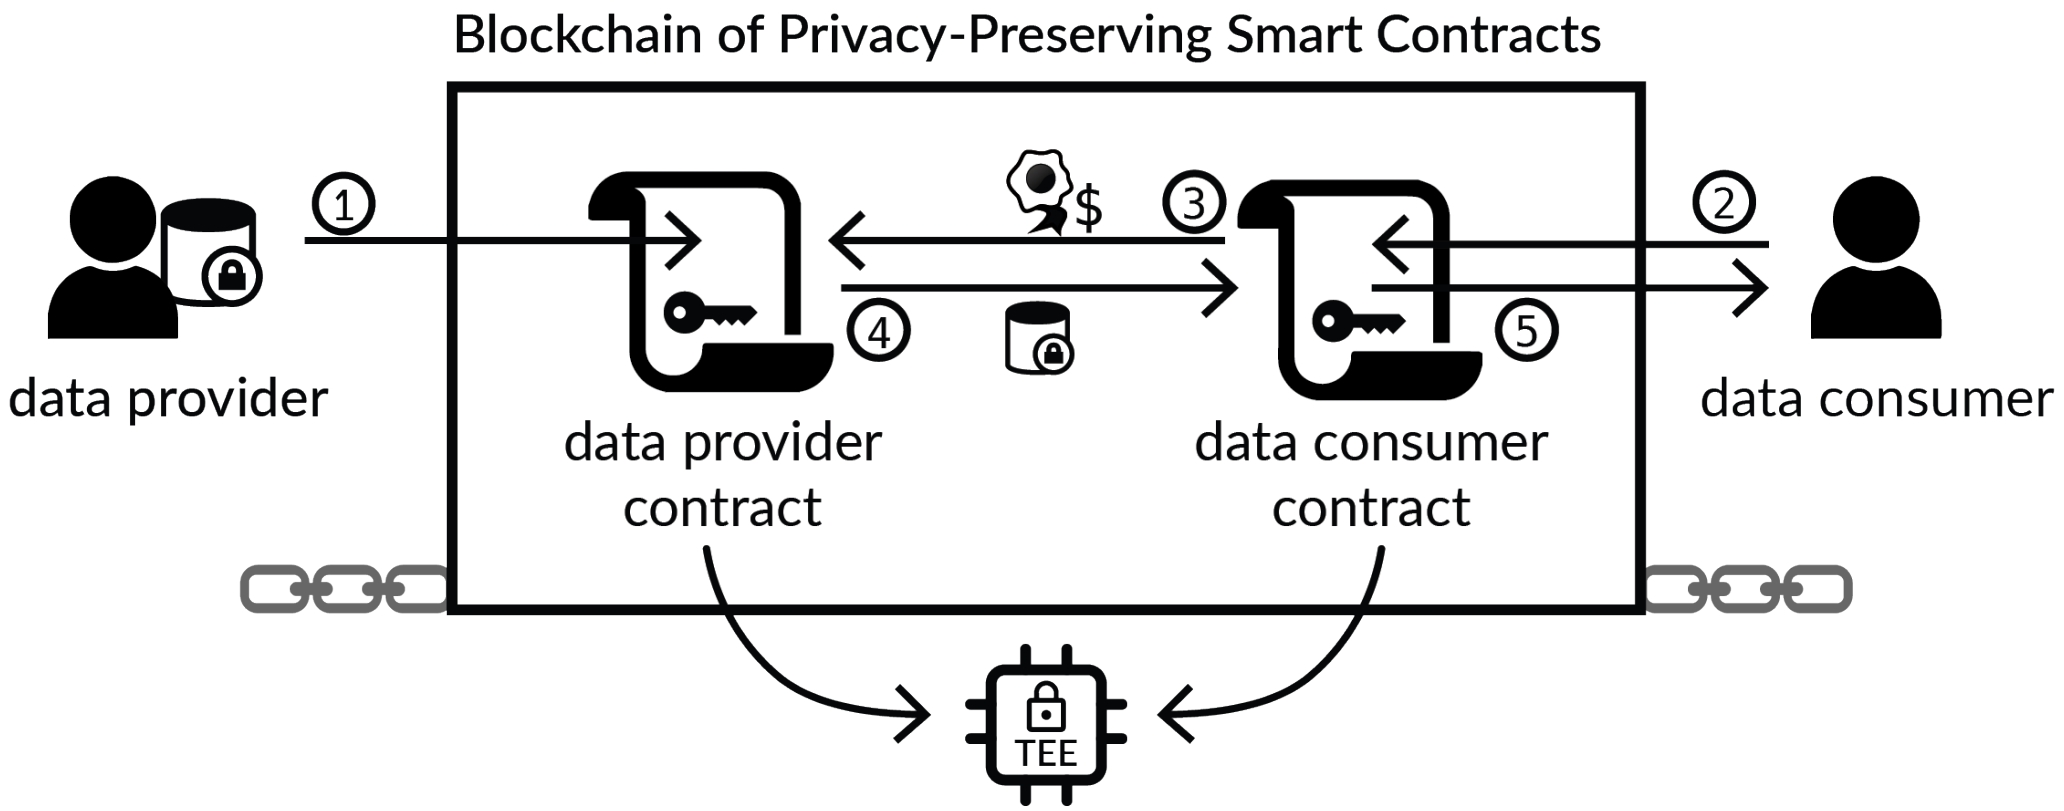

A Demonstration of Sterling: A Privacy-Preserving Data Marketplace. N. Hynes, D. Yan, R. Cheng, and D. Song. VLDB 2018.