

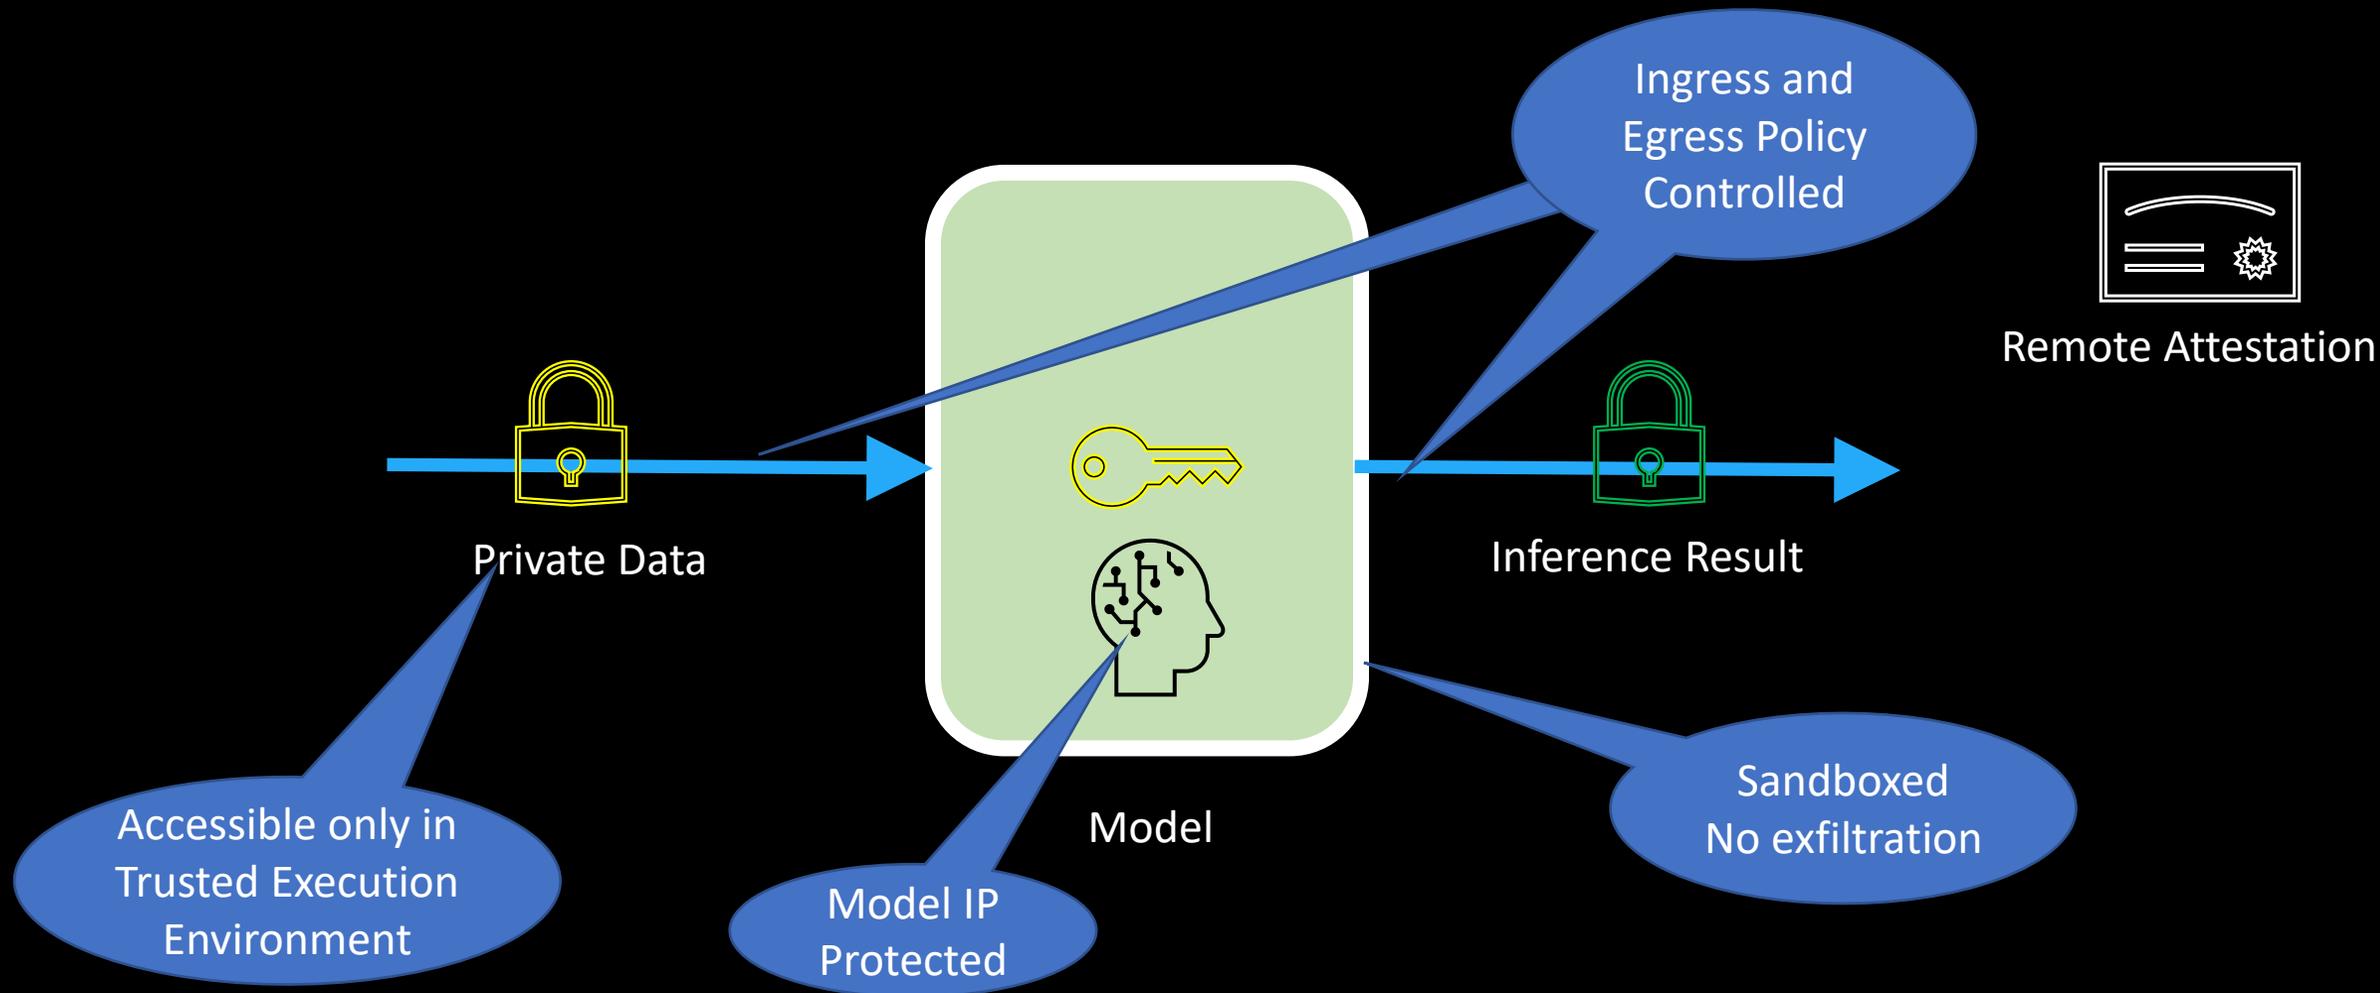
Microsoft Research

Summit 2022

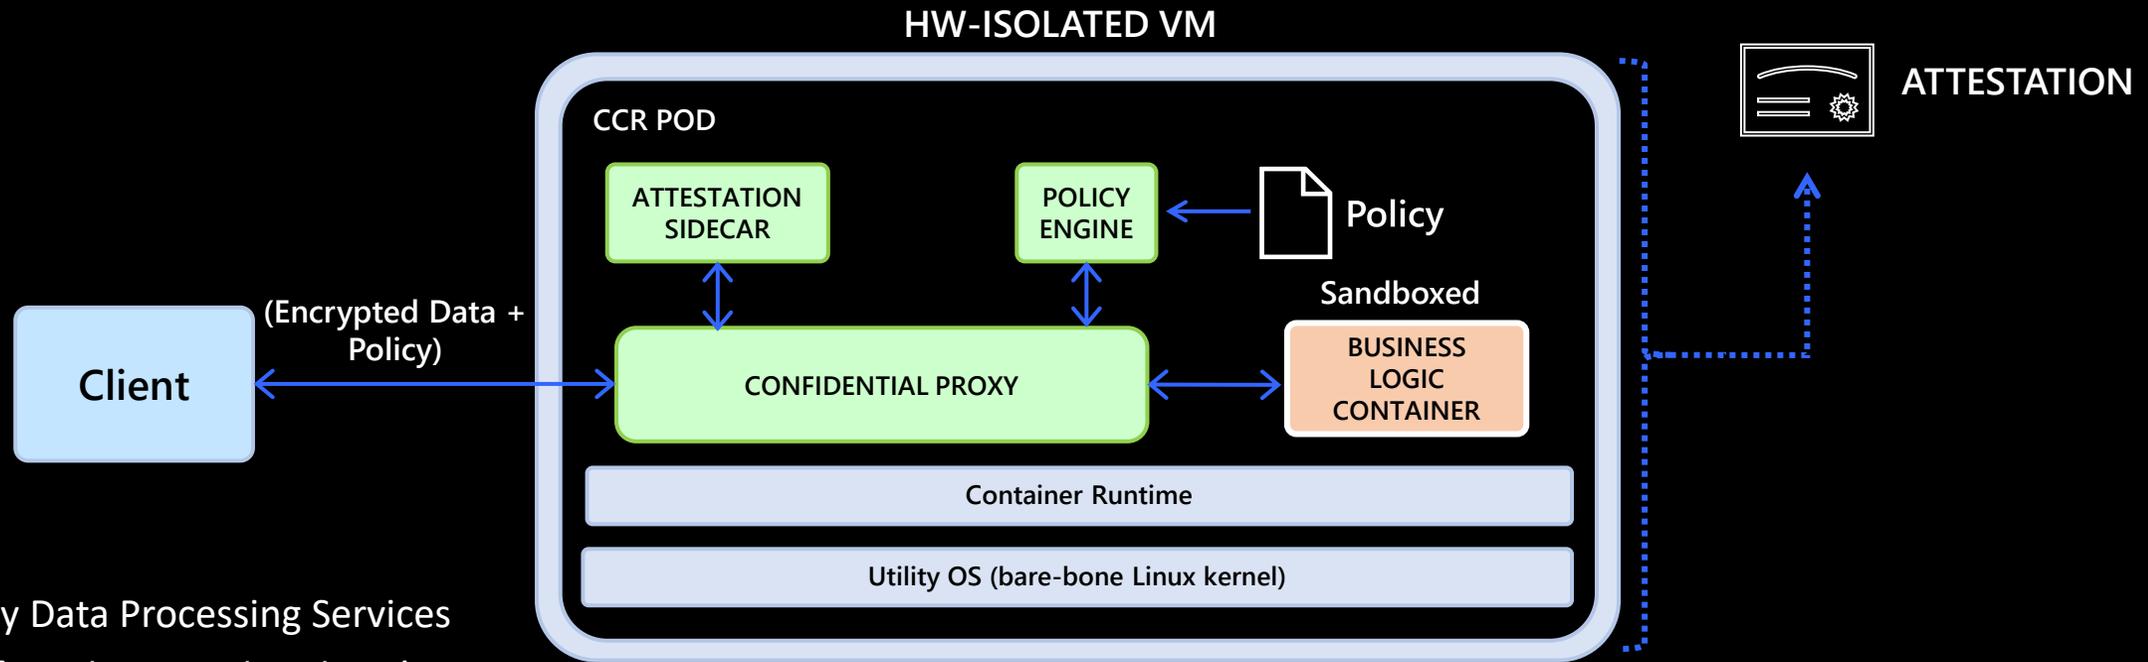
Confidential Clean Rooms for Compliant AI

Satya Lokam (Microsoft Research India)

Privacy Preserving Inference



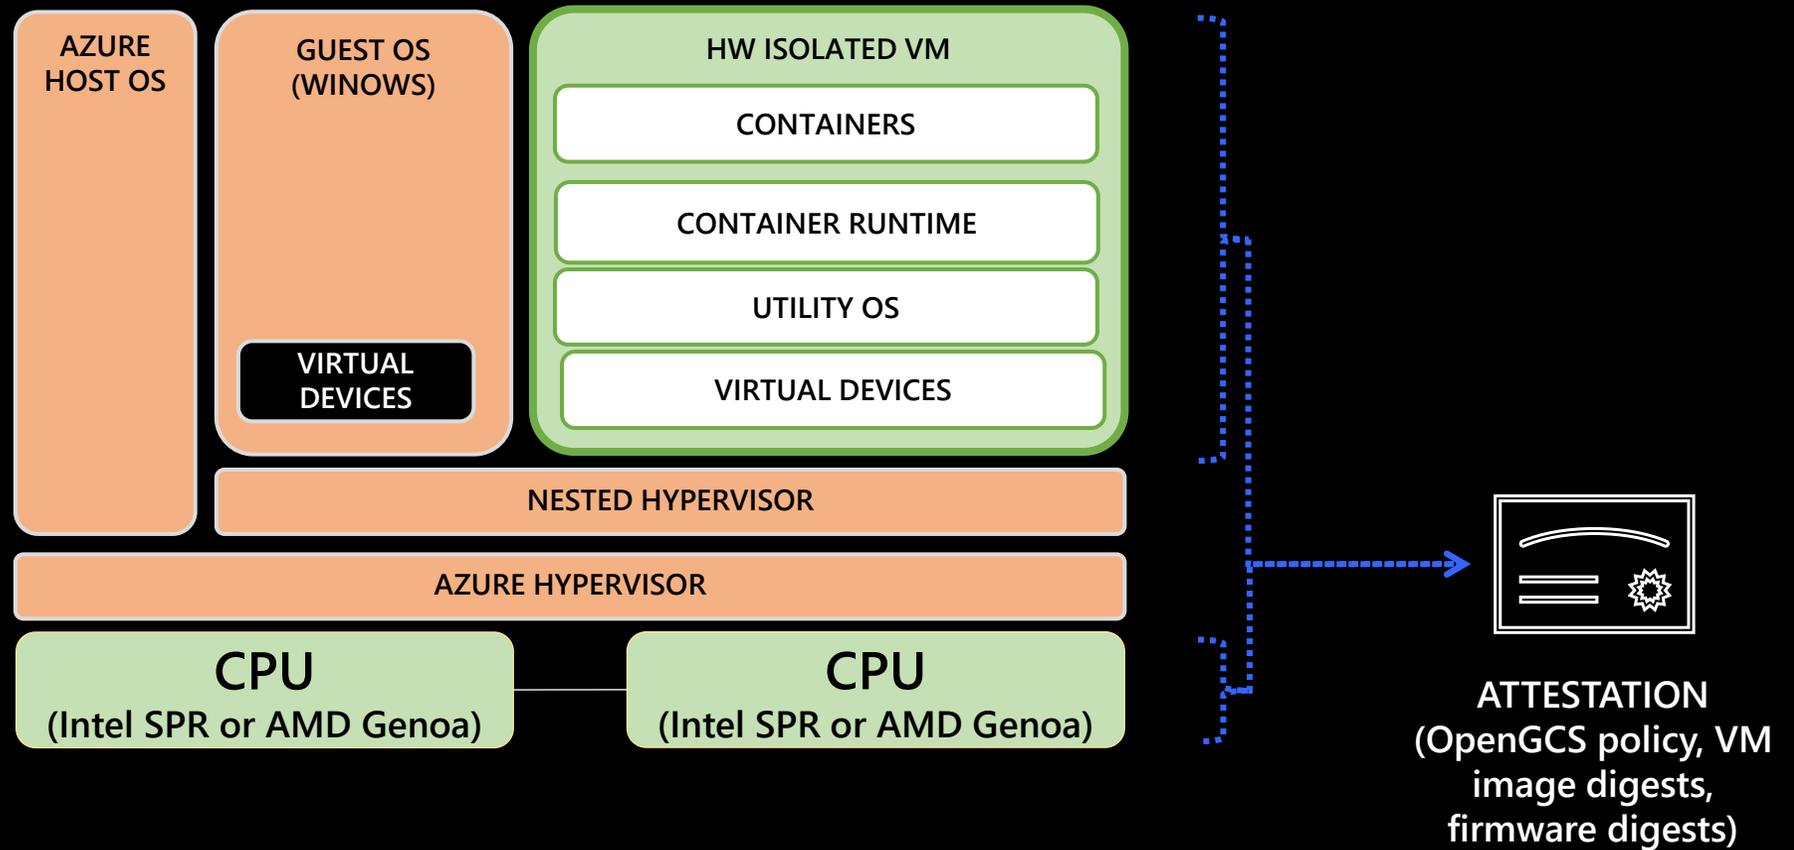
Confidential Clean Rooms



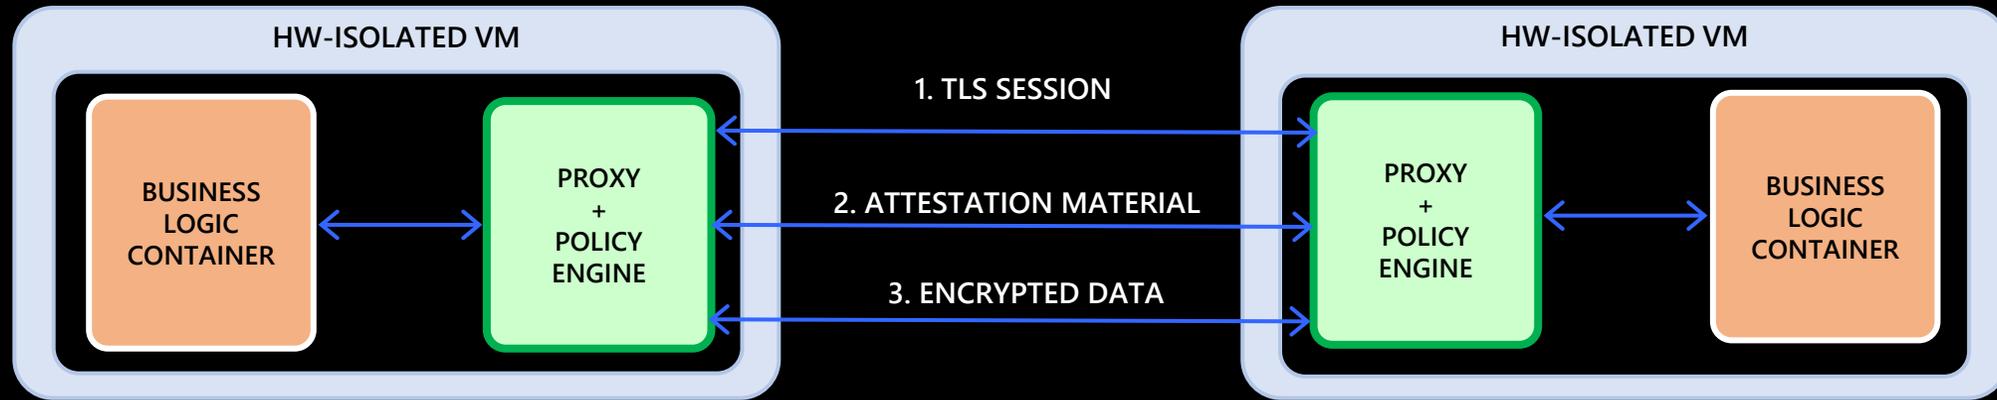
CCR Pod Architecture

- Trustworthy Data Processing Services
- Built out of Hardware Isolated VM's , e.g., AMD SEV-SNP
- Microservices hosted in *Pods* -- CCR
- **Untrusted** Business Logic runs in sandboxed containers, e.g., **ONNX code of an ML model**
- **Trusted** sidecar containers run attestation of TEE's, configs, etc. , do policy checks, . . .
- Ingress and egress policy checks (OPA + Rego) on data intercepted by a trusted proxy (Envoy)

Confidential ACI (Azure Container Instances)



Mutually Attested TLS for composition

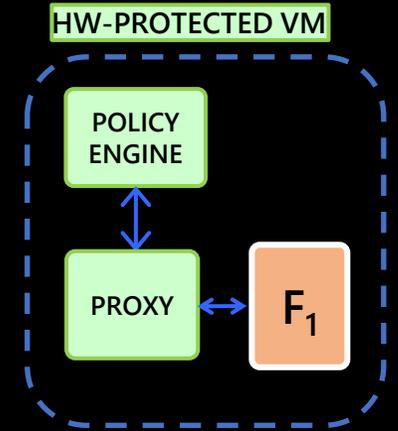


- A standard TLS session is established between sidecar proxies
- Attestation material is exchanged over the TLS session
- Peer's attestation material is verified by the proxy
- If the verification succeeds, data is transferred using standard TLS encryption
- Otherwise, the connection is closed

What do we get?

- Data decrypted only inside the CCR
 - *Confidentiality* (even from platform)
- Attestation
 - Remote *verification before* processing
- Proxy + Policy checks
 - All *ingress/egress* data intercepted and *checked for compliance*
- Sandboxed business logic container
 - *No exfiltration*
 - *IP protected*
 - *“Lift-and-Shift”*
- Attested TLS
 - *Composability* of confidential microservices
- Envoy proxy, OPA + Rego for policy, Conf. Comp
 - Friendly to cloud native dev ecosystem

Privacy and Policy Compliance

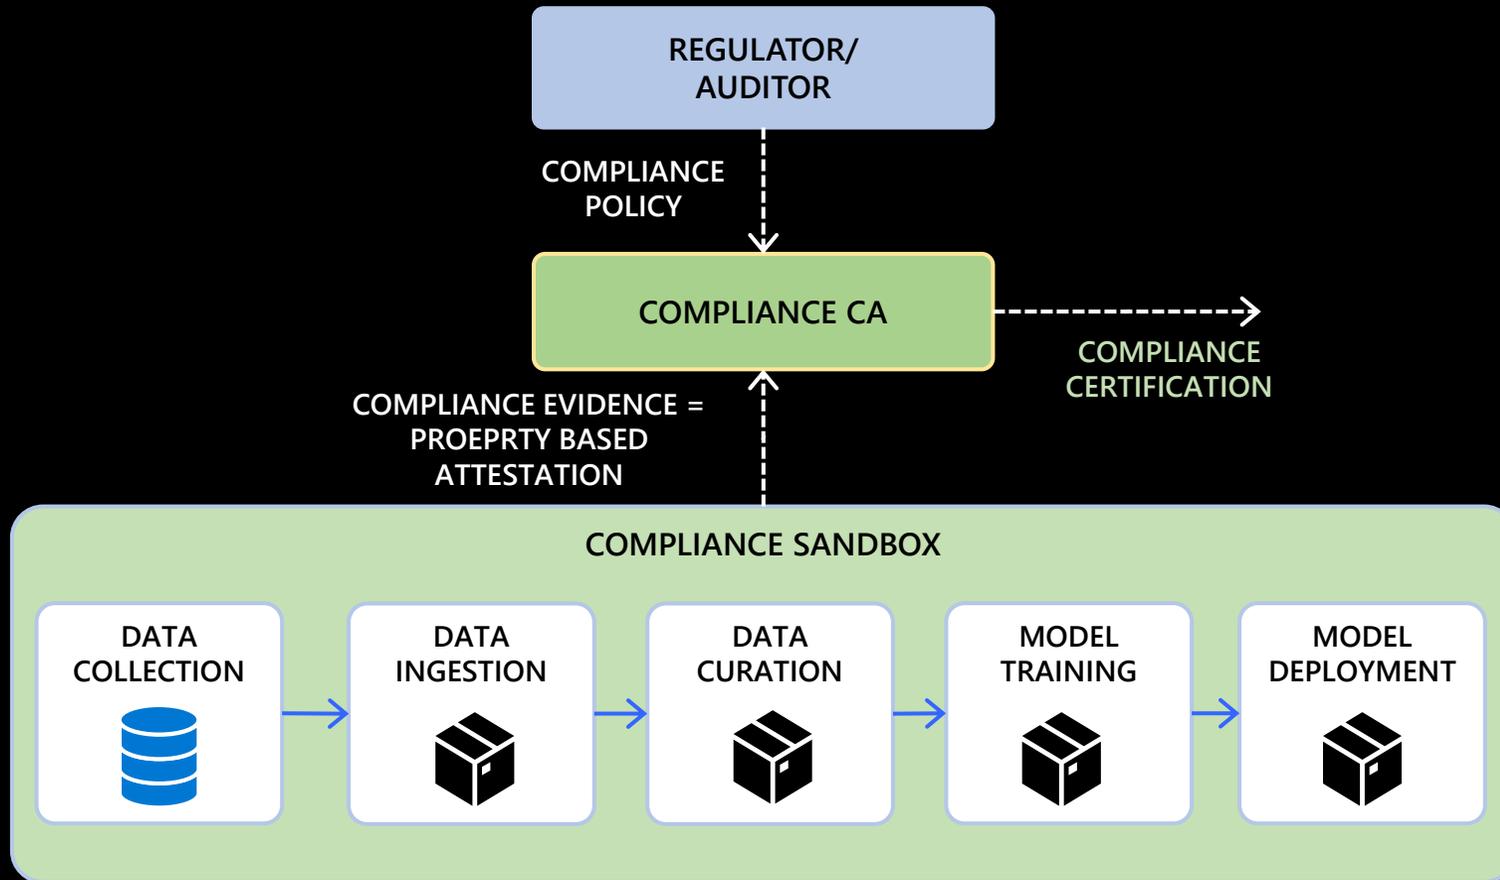


Ease of adoption
Usability



Goal: End-to-end Compliance Framework

- Automated
- Online
- Low-trust
- Multi-tenant
- Protects model IP



Progress so far



View confidential clean rooms as
next evolution of DEPA
(Data Empowerment and Protection Architecture)



Pilot on *Cash-Flow Lending*

Demonstrate feasibility and value
Multiple startups already on-board

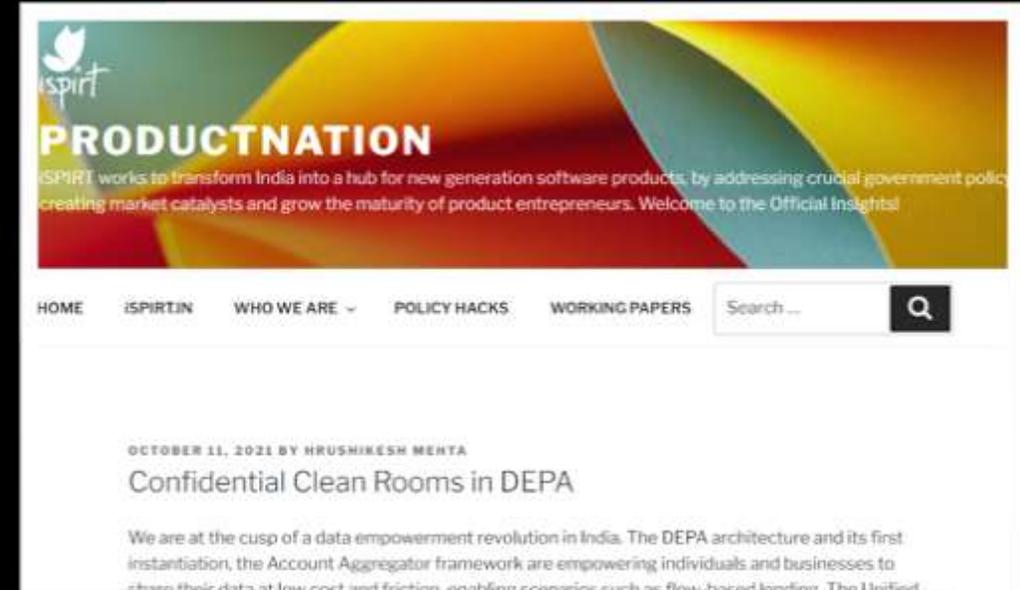


Standards/Guidelines

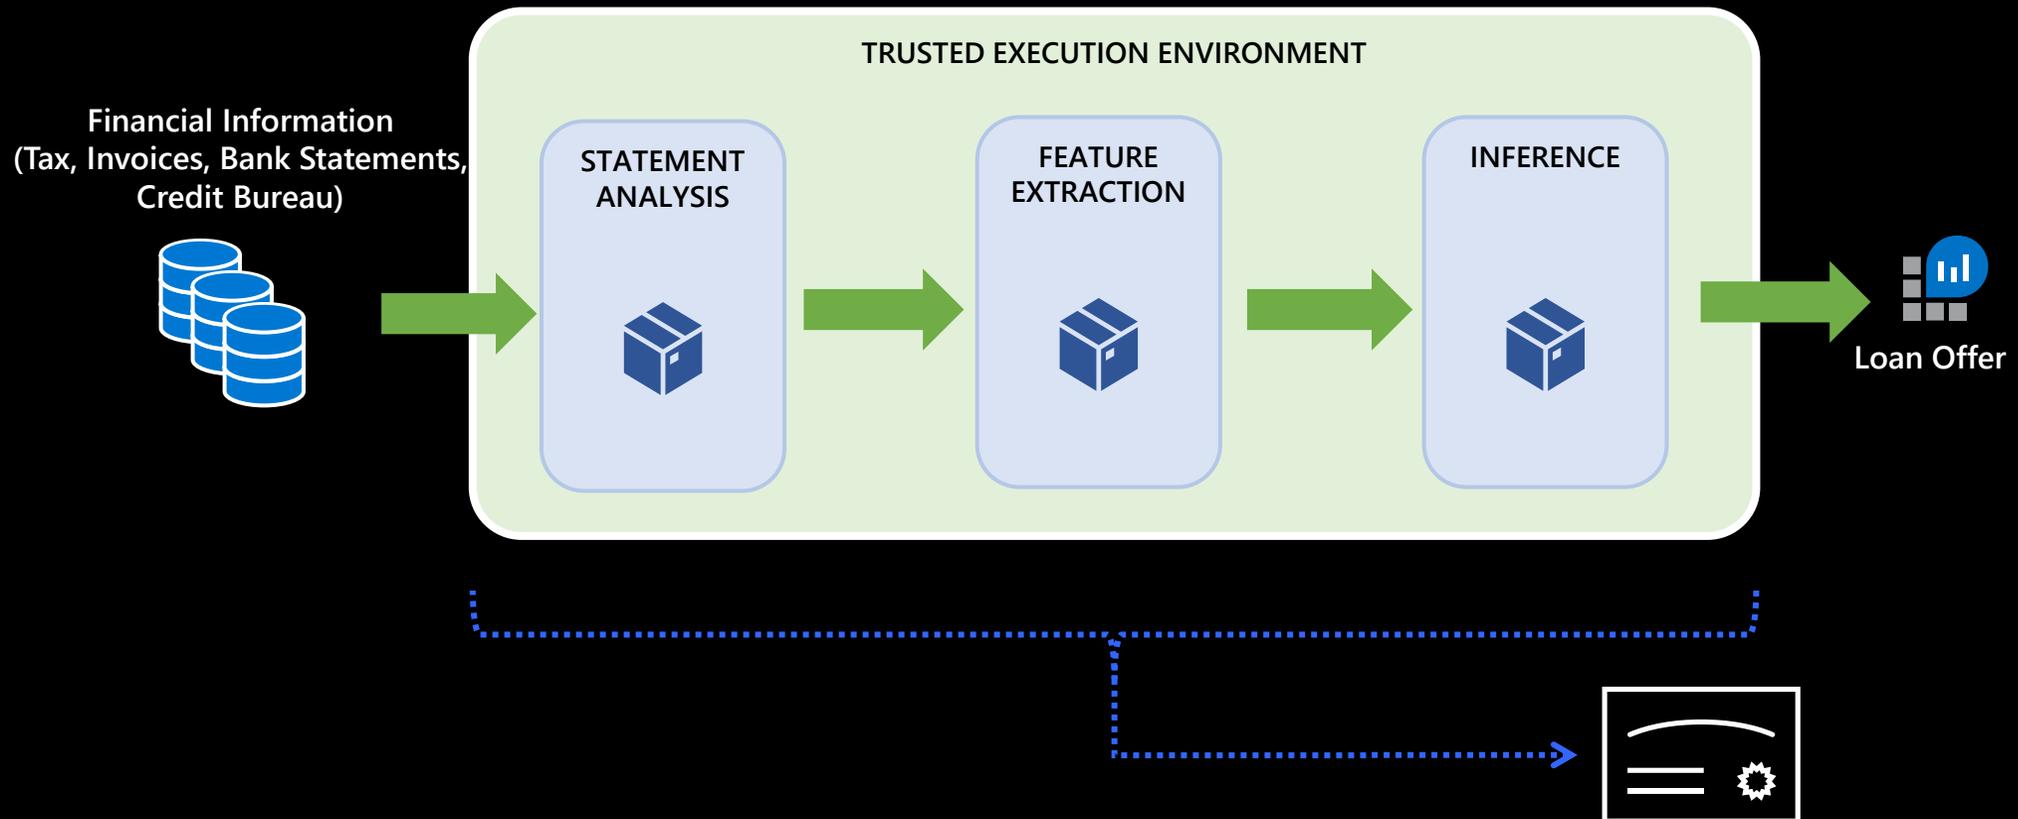
CCR v1.0, Reference implementation



Work in progress!



Example Scenario : Loan Processing



Challenges

- Validating *algorithmic* privacy guarantees of CCR business logic
 - e.g., DP guarantees of a training pipeline,
 - Attribute Inference from a trained model
- Certification + Attestation
 - Trusted Third Party off-line verification and certification
 - Embed that certification in attestation of Confidential Sandboxed Container
- Auditing tools for Privacy of ML models (off-line + certification)
 - Auditing: Bayesian estimation of DP (<https://arxiv.org/pdf/2206.05199.pdf>)
 - Quantifying membership/attribute inference ([Yeom et al.](#))

Microsoft Research
Summit 2022

Thank you

Stay in touch: [@satya](#)