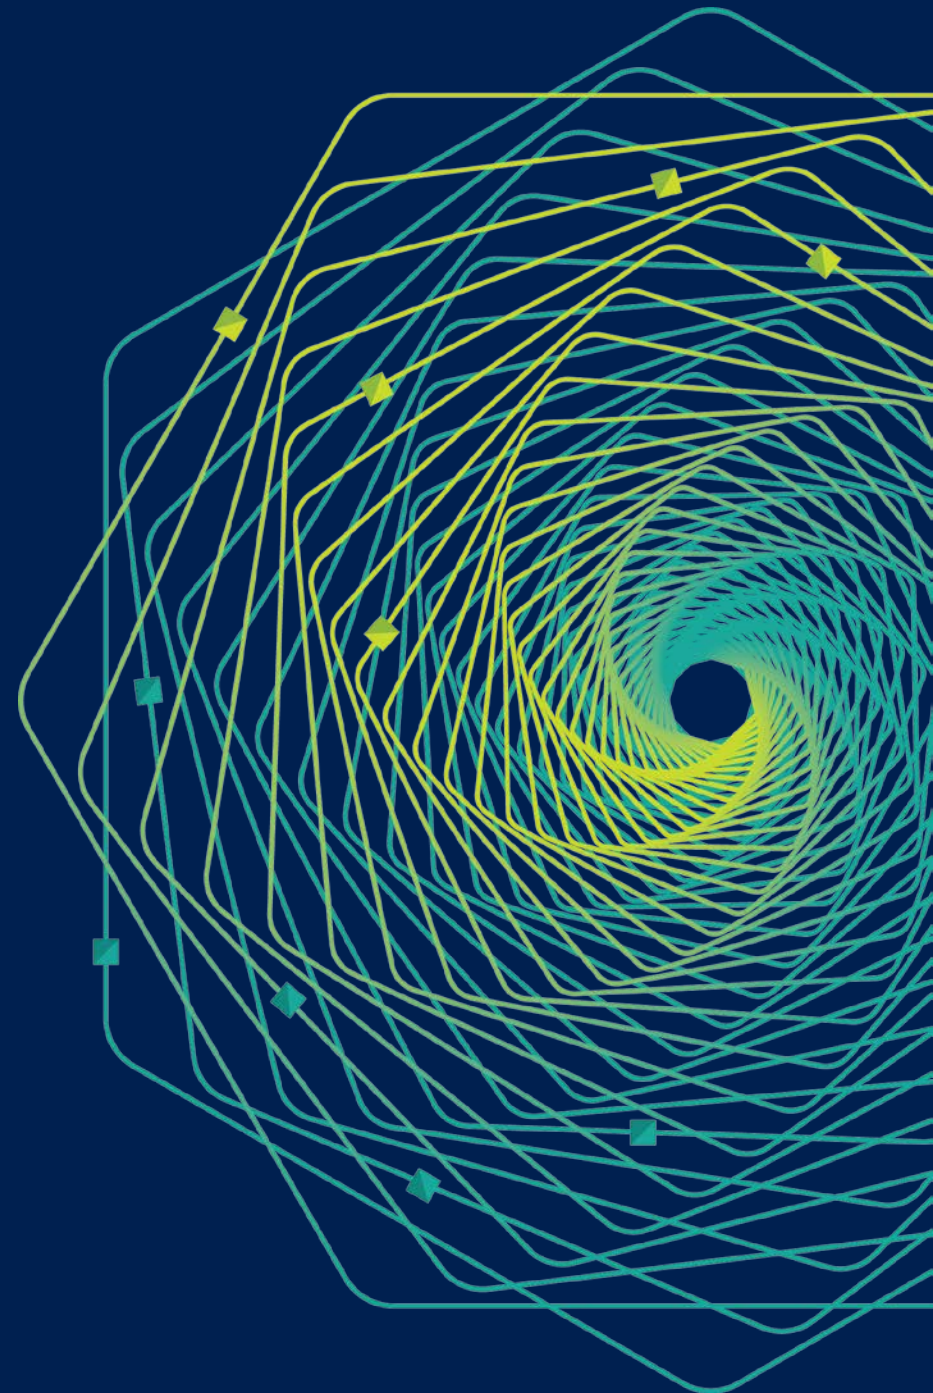



Research Faculty Summit 2018

Systems | Fueling future disruptions



The Rise of Confidential Computing

Mark Russinovich
CTO, Microsoft Azure, Microsoft
[@markrussinovich](https://twitter.com/markrussinovich)



Cloud Data Threats

Customer cloud data concerns:



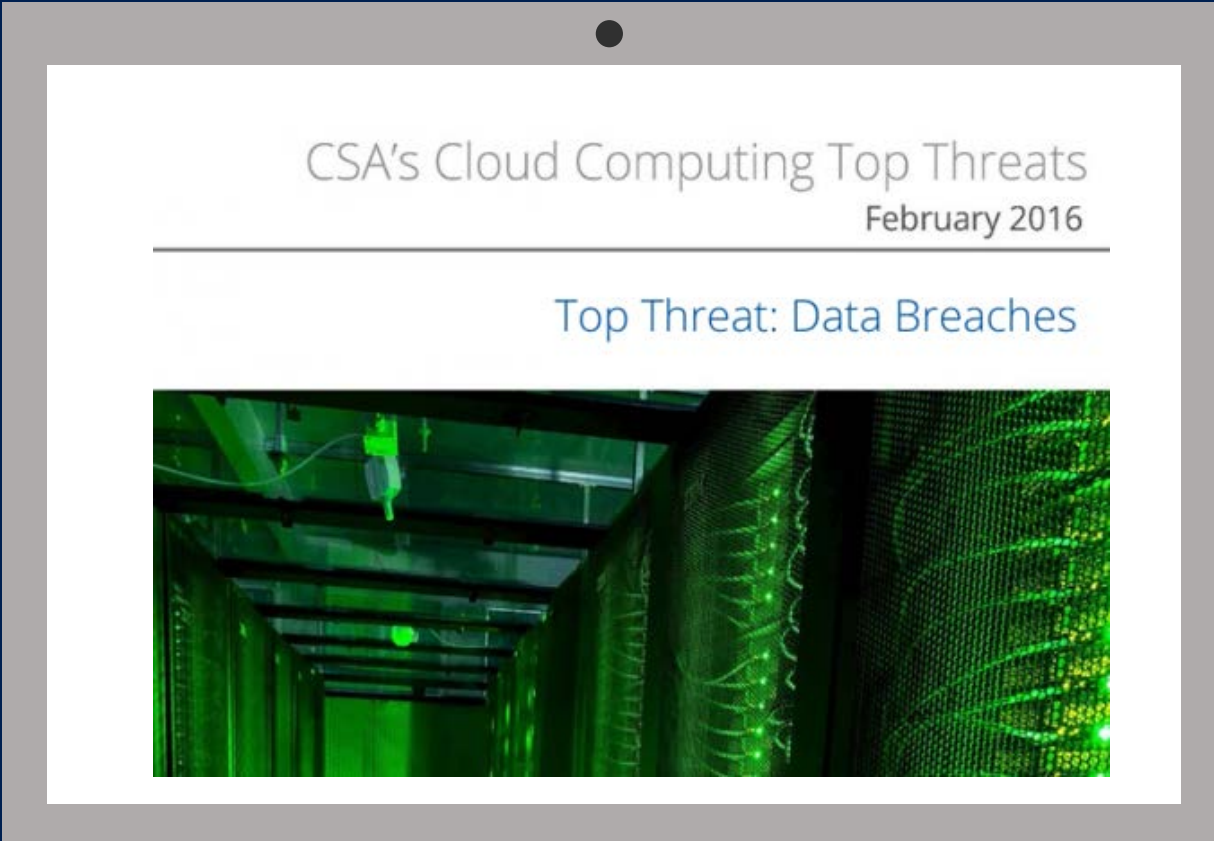
Malicious privileged admins or insiders



Hackers exploiting bugs in the infrastructure



Third-party access without customer consent



Data Protection

At rest



Encrypt inactive data when stored in blob storage, database, etc.

Examples:

- Azure Storage Service Encryption for Data at Rest
- SQL Server Transparent Database Encryption (TDE)

In transit




Encrypt data that is flowing between untrusted public or private networks

Examples:

- HTTPS
- TLS

In use



Protect/Encrypt data that is in use during computation

Examples include:

- Trusted Execution Environments
- Homomorphic encryption

Trusted Execution Environments (TEEs)

Protected container:

- Isolated portion of processor & memory
- Code & data cannot be viewed or modified from outside

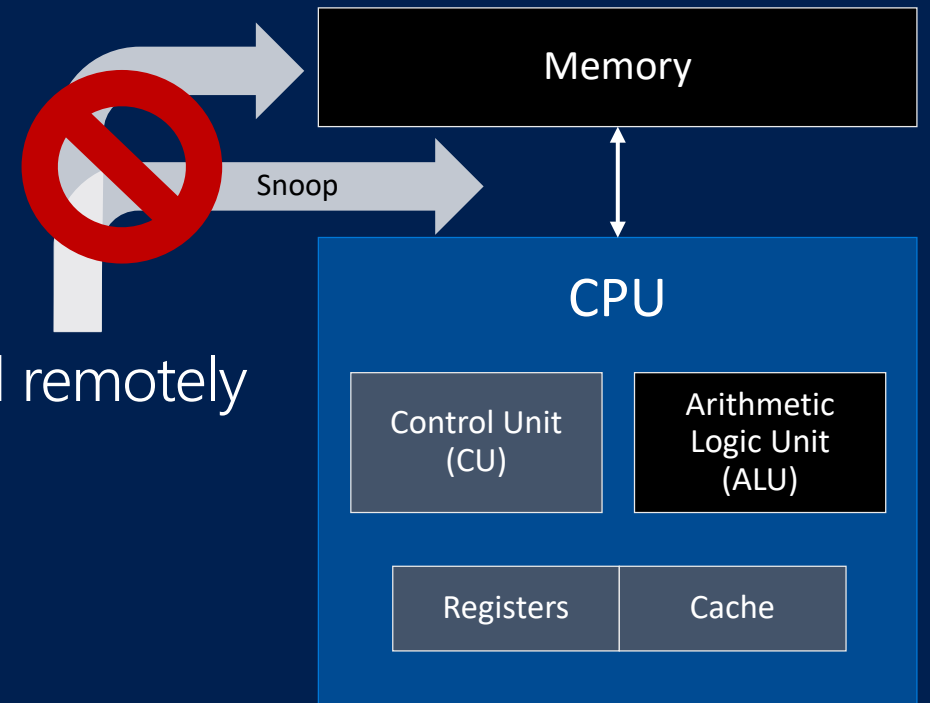
Supports attestation: proving of identity both locally and remotely

Supports sealing: persisting secrets

Examples:

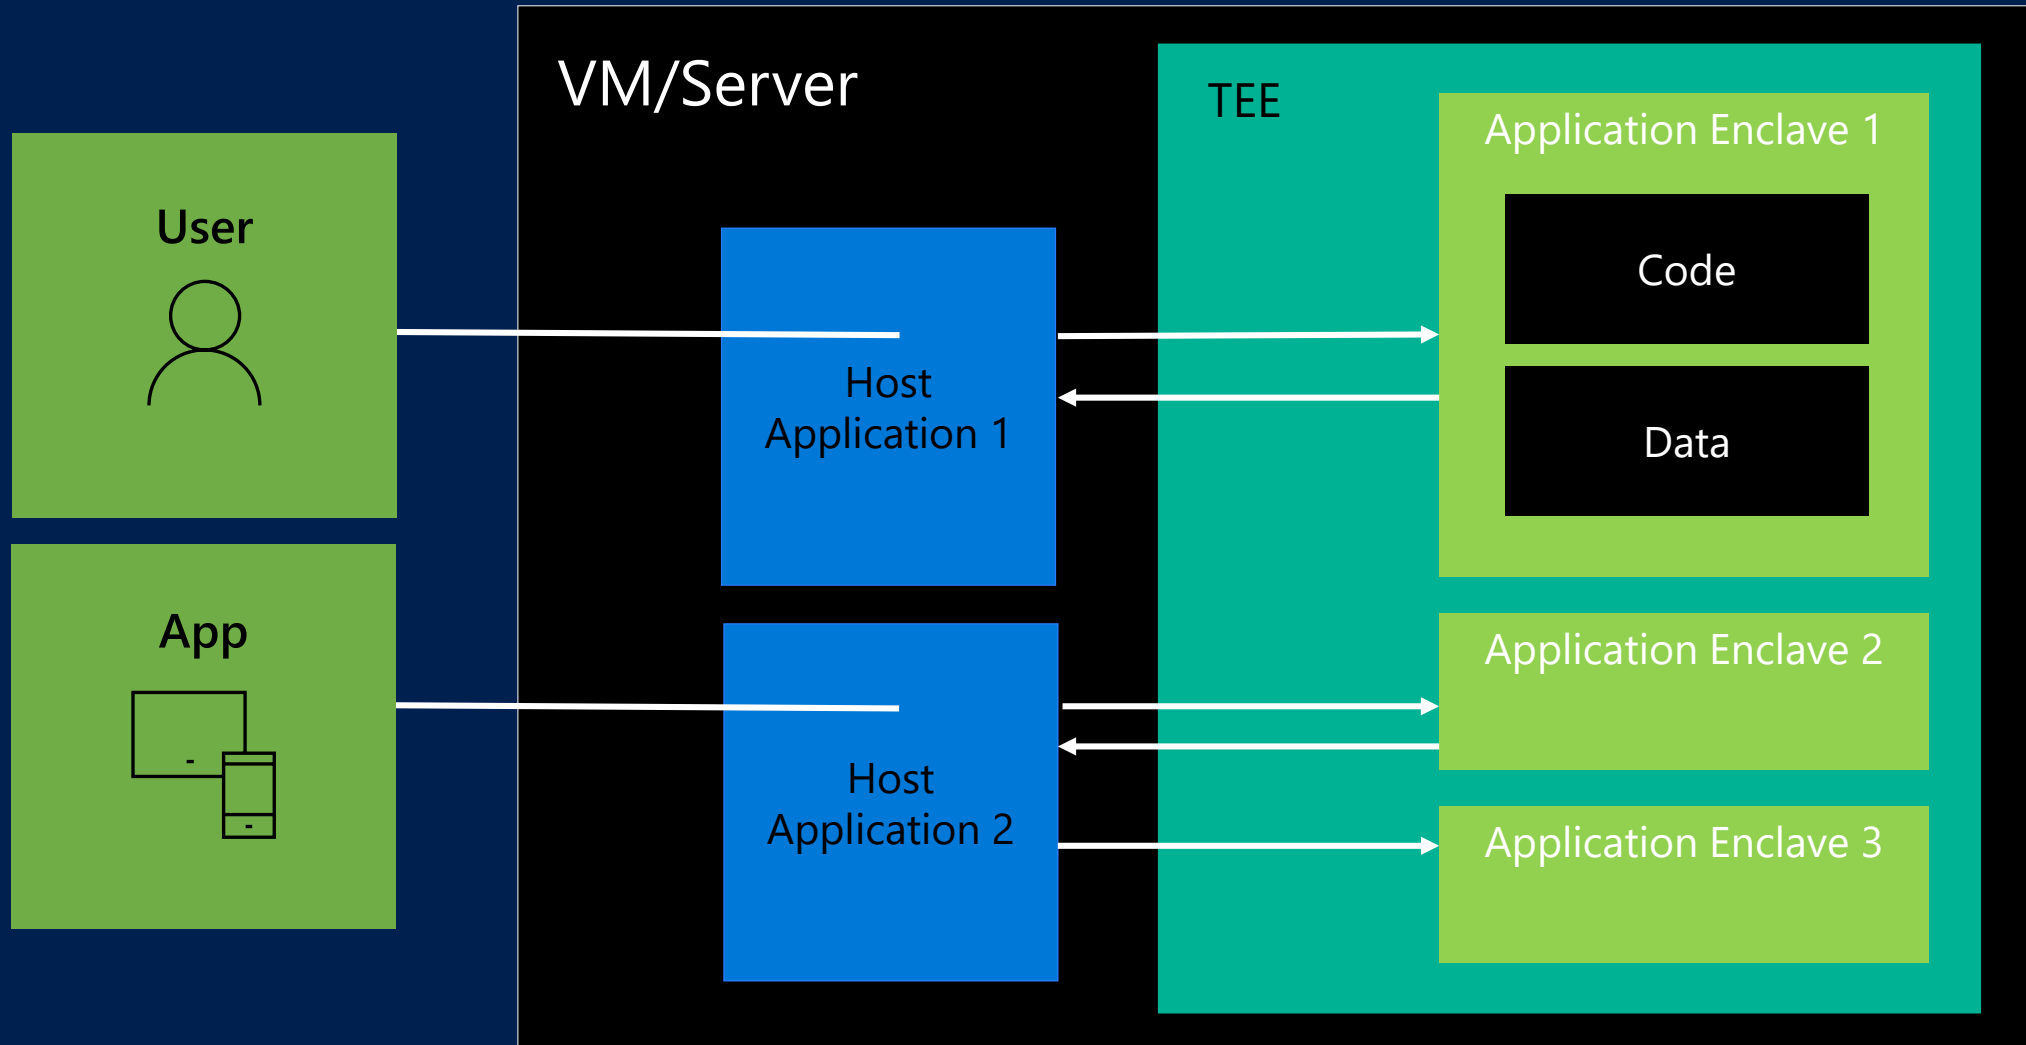
Intel SGX

Virtualization Based Security (VBS) aka Virtual Secure Mode



Hardware-based TEE

TEE application architecture



Azure Confidential Computing

Azure and confidential computing



Working with silicon partners to enable Confidential Computing

Building software to deploy, manage, and develop secure TEE applications on Azure

Designing and developing services to support attestation in the cloud

Enabling confidential PaaS and SaaS services

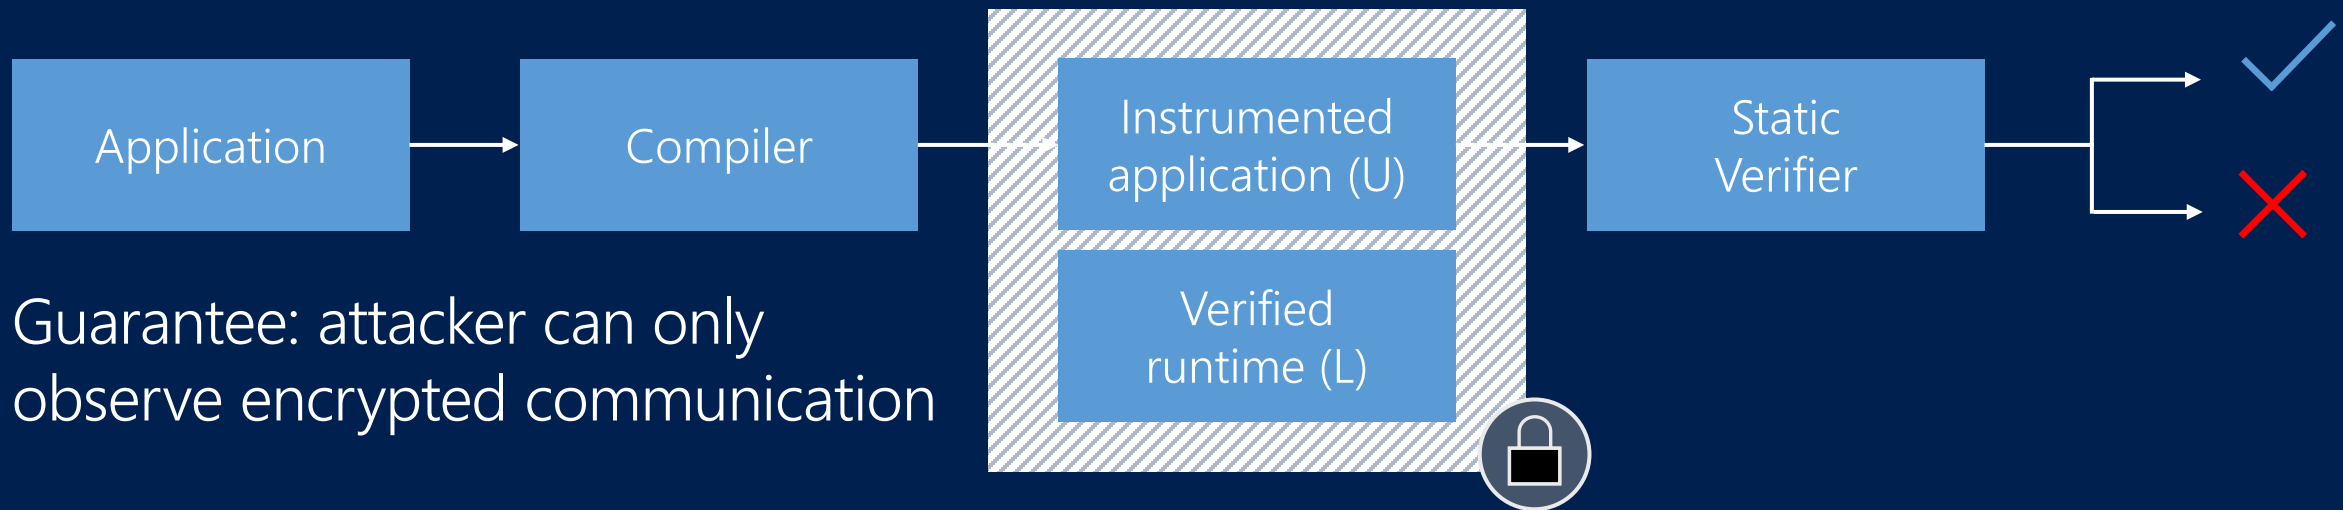
Preventing direct information leaks



Problem: code in enclaves may unintentionally write secrets out



Solution: use a compiler that instruments memory accesses & verify that instrumented binary does not leak secrets



Preventing indirect information leaks

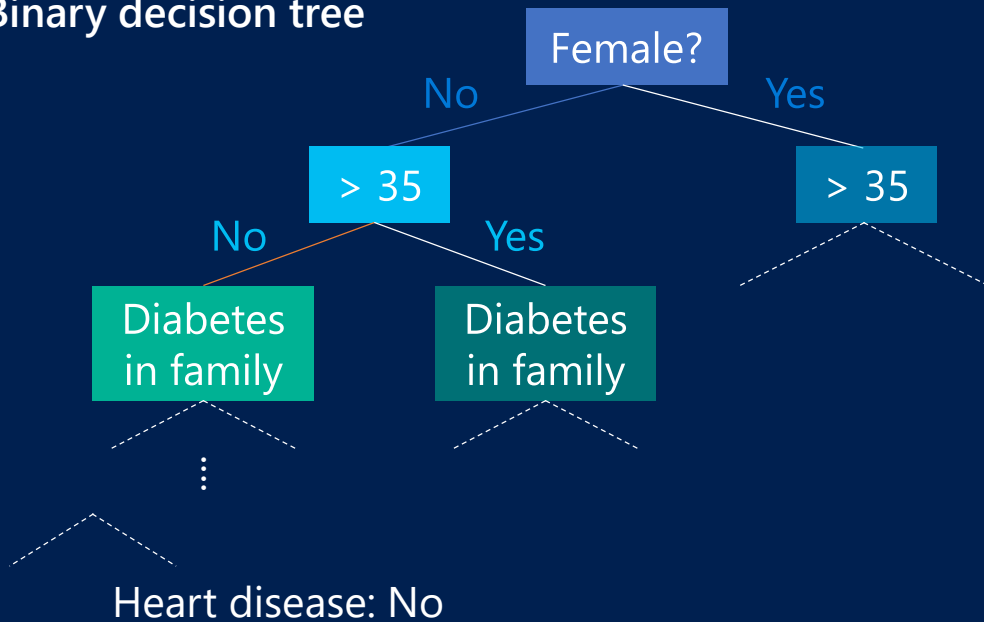


Problem: memory/disk access patterns may leak information



Solution: use compiler and hardened libraries that prevent leaks with data oblivious primitives

Binary decision tree



Memory



Demo: Oblivious computing

Example confidential computing scenarios




Always encrypted storage
with
SQL Server



Enabling scalable
and confidential
blockchain networks
with Coco Framework



Financial data
processing

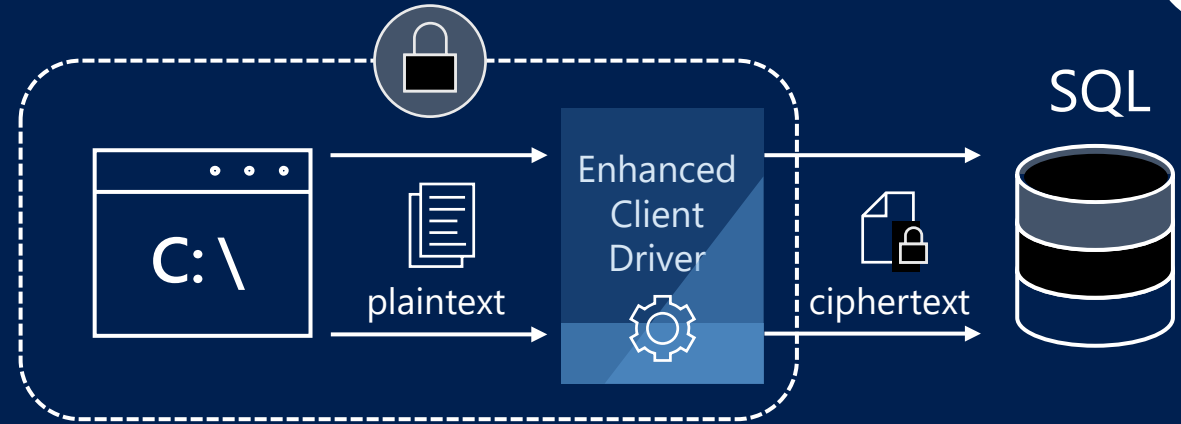


Secure multi-party
machine learning

SQL Always Encrypted

Protects sensitive data in use from high-privileged yet unauthorized SQL users both on-premises and in the cloud

Current GA version in SQL Server 2016/17 and Azure SQL DB



Client side Encryption

Client-side encryption of sensitive data using keys that are *never* given to the database system

Encryption Transparency

Client driver transparently encrypts query parameters and decrypts encrypted results

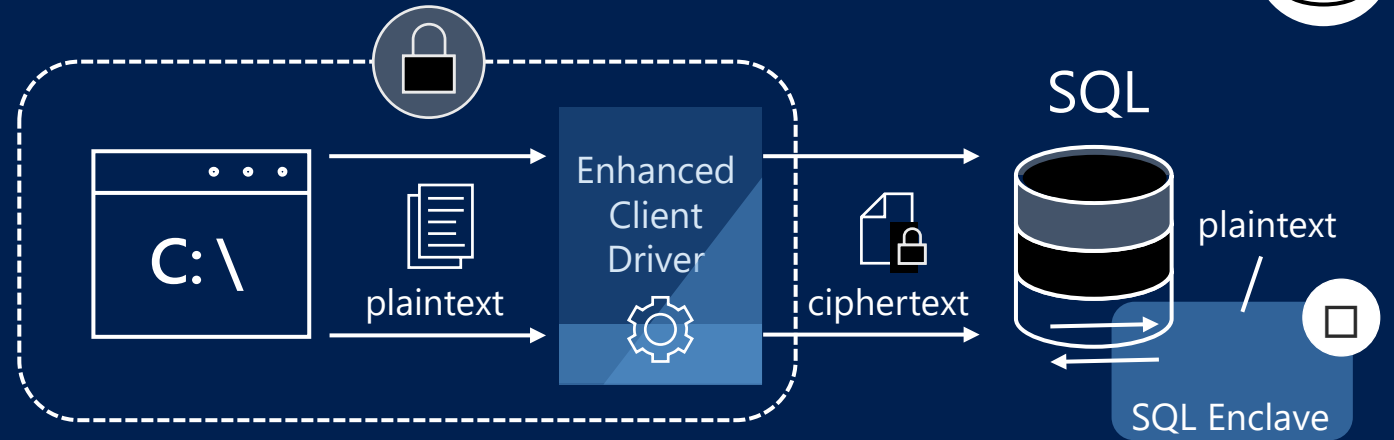
Queries on Encrypted Data

Support for equality comparison, including join, group by and distinct operators via deterministic encryption

Confidential SQL Always Encrypted



Protects sensitive data in use **while preserving rich queries and providing in-place encryption**



Secure computations inside SQL Enclave

SQL Server Engine delegates operations on encrypted to the SQL Enclave, where the data can be safely decrypted and processed

Rich Queries

pattern matching (LIKE), range queries (<, >, etc.), sorting, type conversions, support for non-bin2 collation, and more

In-place Encryption

SQL Enclave can perform initial data encryption and key rotation, without moving the data out of the database

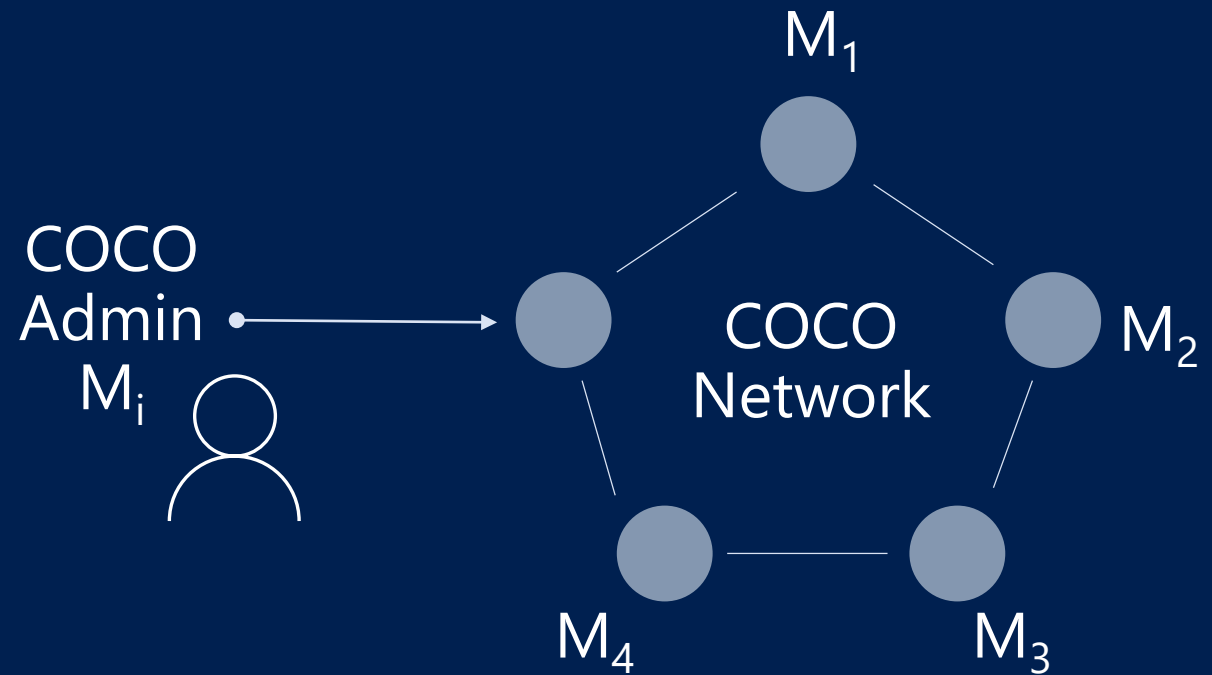
Coco Framework: Confidential Consortium Blockchain Framework



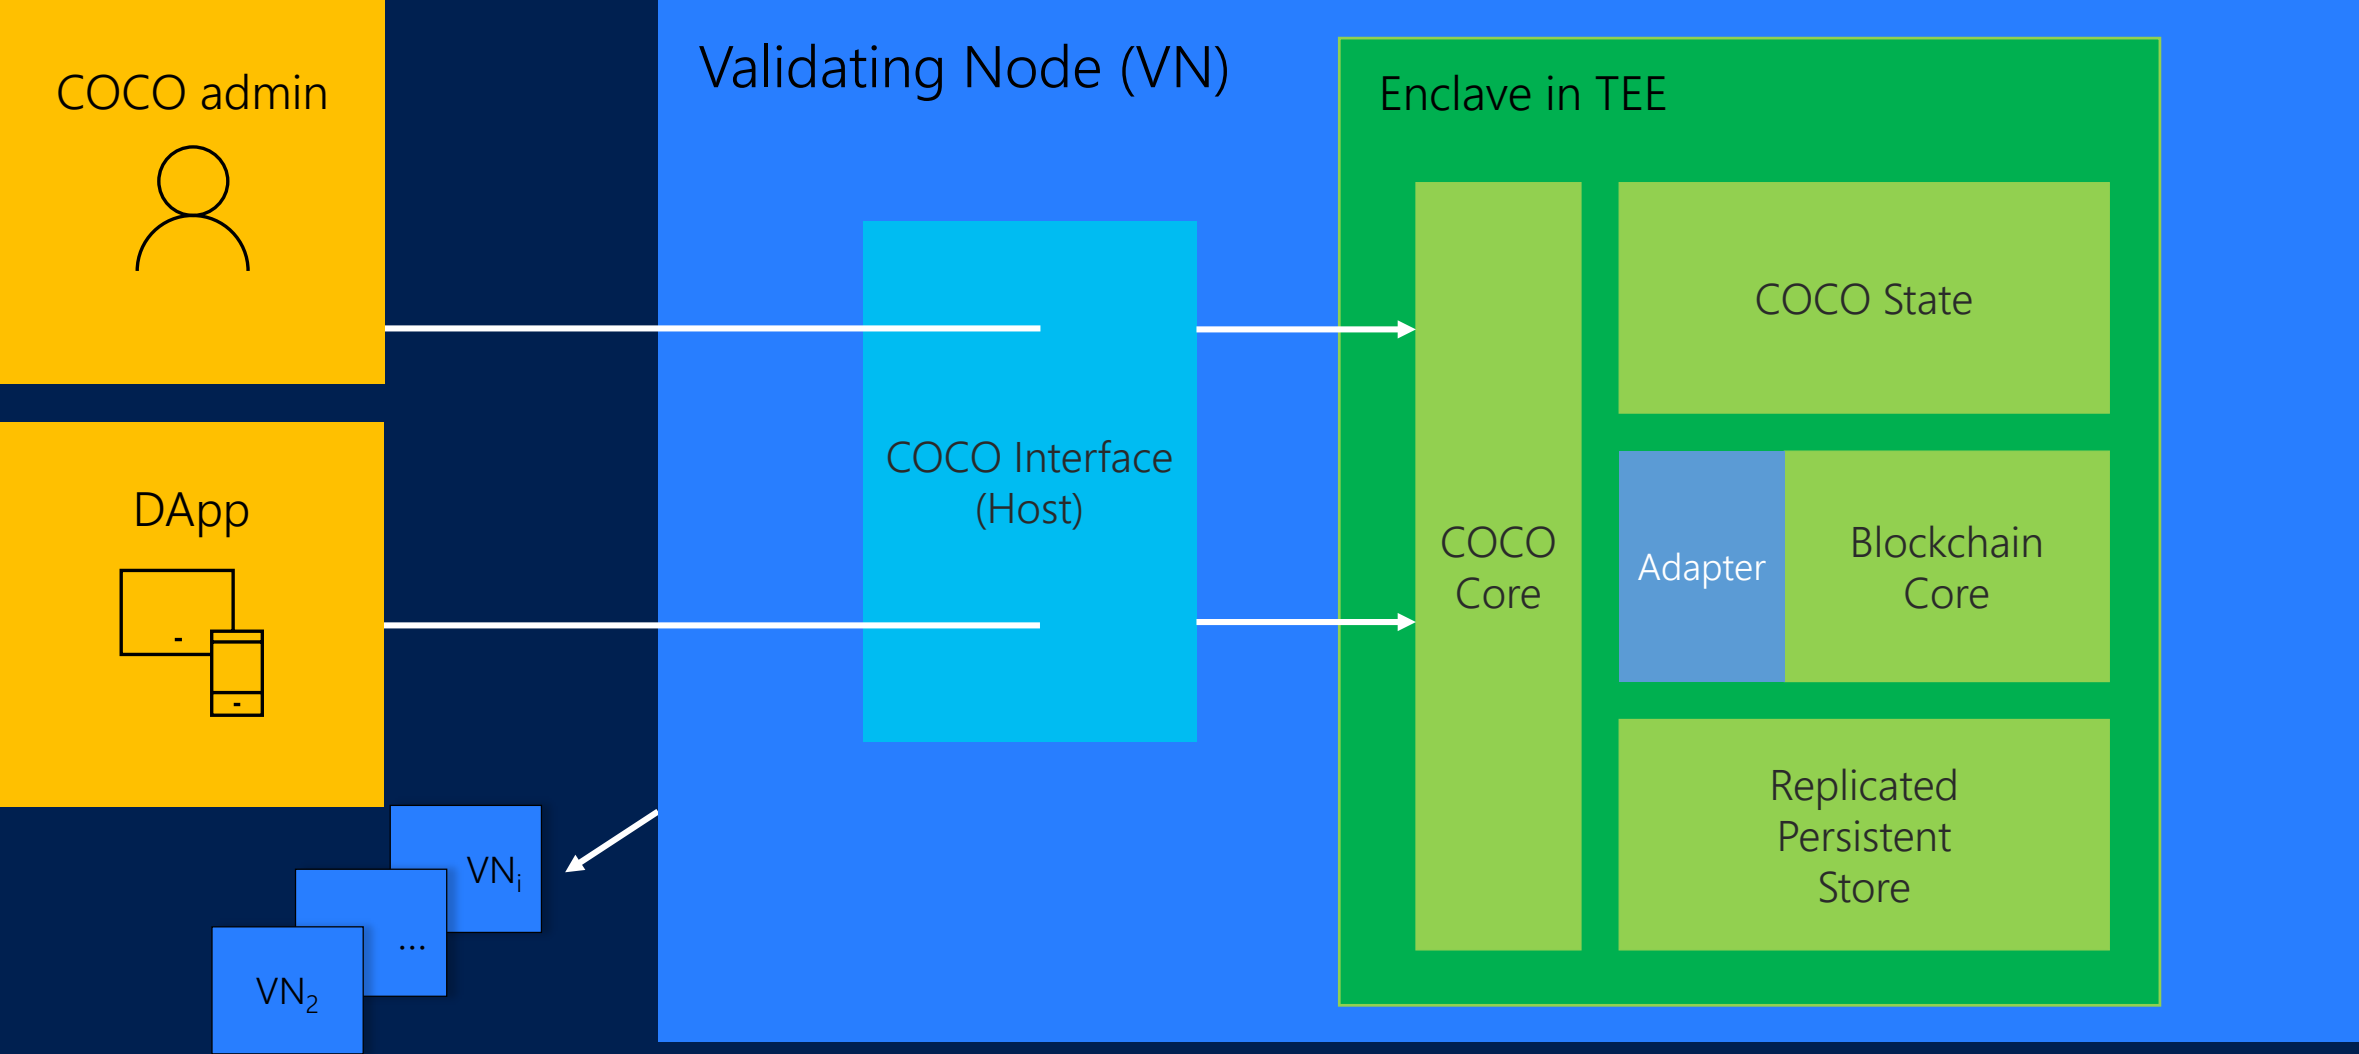
Open-source framework that enables high-throughput (~100x), fine-grained confidentiality, and consortium governance for blockchain

Creates a trusted network of physical nodes on which to run a distributed ledger, providing secure, reliable components for the protocol to use

Through the use of TEEs able to simplify consensus and transaction processing



Coco Framework architecture



Demo:

Coco Ethereum versus Ethereum

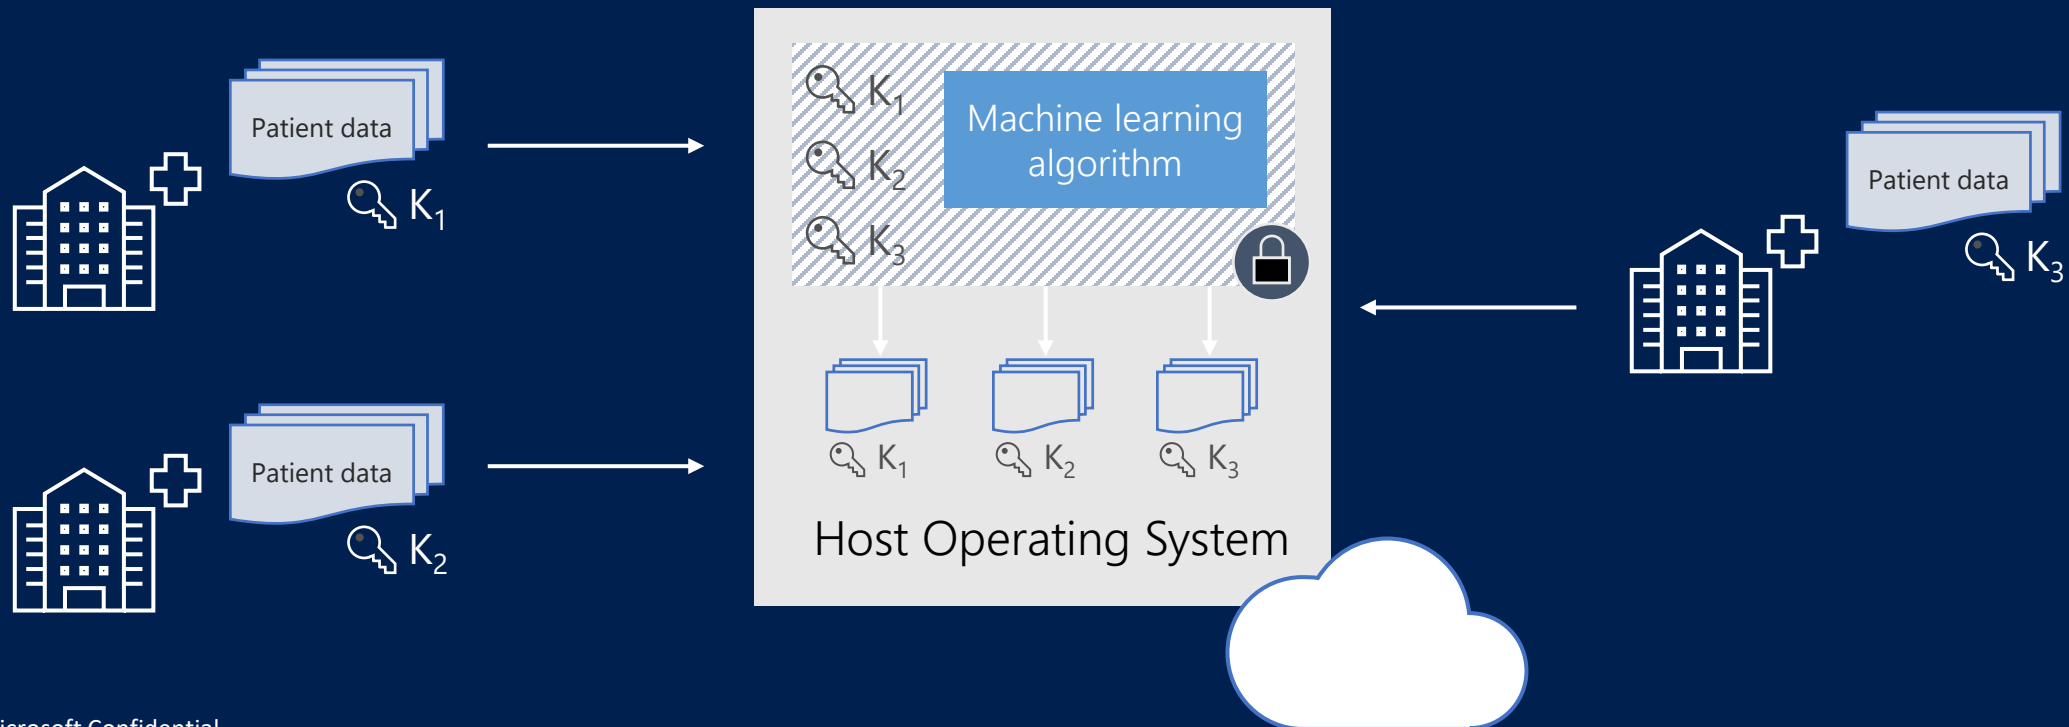
Confidential multi-party machine learning



Partnered health facilities contribute private patient health data sets to train a ML model

Each facility only sees their respective data sets (aka no one, not even cloud provider, can see all data or trained model, if necessary)

All facilities benefit from using trained model




Demo:

Confidential multi-party ML


Summary



Confidential computing
in the cloud is in its
early stages



Microsoft is driving the
direction & adoption of
newer trusted execution
environments in the cloud



Azure is empowering
new secure business
scenarios in the cloud

References

Blockchain with Coco Fx:
<http://aka.ms/cocopaper>

Multi-party machine learning:
<https://www.microsoft.com/en-us/research/wp-content/uploads/2016/07/paper.pdf>

SQL Server with Haven:
<https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/osdi2014-haven.pdf>

Map/reduce with VC3:
<https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/vc3-oakland2015.pdf>

Preventing enclave information leaks:
<https://people.eecs.berkeley.edu/~rsinha/research/pubs/pldi2016.pdf>

Using side-channel page faults to extract JPG images:
<https://www.microsoft.com/en-us/research/wp-content/uploads/2017/06/atc17-final230.pdf>

Thank you!

