EscrowAI System Capabilities

Overview

EscrowAI® is a patented zero-trust, confidential computing platform for secure collaboration between algorithm developers and controllers of protected data. With EscrowAI, data remains within the data controller's secure environment and is made available for computation in a hardware-based Trusted Execution Environment (TEE) instance.

EscrowAI is protected by U.S. Patents 11,531,904; 11,748,633; 12,001,965; 12,093,423; 12,099,630; 12,111,951, 12,141,319, 12,229,274, 12,339,993 and other patents pending.

Version: 2.1.0

Security

EscrowAI has the following important security features:

- **Data Controllers retain control of their data.** The Data Controllers' organizations retain full control of the data. With EscrowAI, data stays in the Data Controllers' organizational infrastructure.
- **Strong encryption protects data at rest, in transit, and in use.**
  Data is encrypted with AES-256-based ciphers and all network connections use TLS 1.2.
- **Data is protected in use.**
  Data and algorithms achieve high isolation and memory encryption using hardware-based assurances. Secure attestation ensures the authenticity and integrity of the TEE execution environments, and PCR validation ensures the validity of all software and other compute artifacts.
- **Algorithm intellectual property is protected.**
  Algorithm owners retain protection of their intellectual property. They encrypt their models, which are decrypted only in a TEE instance.

Confidential Computing Technology

EscrowAI offers the following types of confidential computing technologies. These technologies underlie the virtual machines (VMs) and other secure resources in the TEE instance that EscrowAI manages for you.

Confidential Virtual Machines

Confidential VMs are a form of TEE that provides elevated protection for customer data from the underlying infrastructure, cloud operators, and BeeKeeperAI. Unlike other approaches and

solutions, you don't have to adapt your existing workloads to fit the platform's technical needs. EscrowAI makes available AMD-based Confidential VMs that use AMD SEV-SNP, which was introduced with 3rd Gen AMD EPYC™ processors.

Secure Encrypted Virtualization-Secure Nested Paging is a TEE technology that offers multiple protections, including memory encryption, unique CPU keys, encryption for the processor register state, integrity protection, firmware rollback prevention, side-channel hardening, and restrictions on interrupt and exceptions behavior. CVM options available in EscrowAI are Azure general purpose (**DCasv5-series and DCadsv5-series**) and memory-optimized (**ECasv5-series and ECadsv5-series**) machine sizes.

VM-isolated Confidential Containers on Azure Container Instances (CC ACI)

Confidential containers on Azure Container Instances are deployed in a container group with a Hyper-Visor isolated TEE that includes a memory encryption key generated and managed by an AMD SEV-SNP capable processor. Data in use in memory is encrypted with this key to help protect against data replay, corruption, remapping, and aliasing-based attacks.

Confidential Virtual Machines with Confidential GPU

This extension of confidential VMs includes the NVIDIA H100 NVL GPU. The Trusted Execution Environment (TEE) spans confidential VM on the CPU and attached GPU, enabling secure offload of data, models, and computation to the GPU. 4th-generation AMD EPYC™ Genoa processors and the NVIDIA H100 Tensor Core GPU power this VM option. One SKU is available, comprised of 40 EPYC cores, 320 GM of memory, and 1 GPU (**NCCads H100 v5 series**). Availability may be limited by Azure region.

Available Functionality

AI/ML training

EscrowAI incorporates MLflow functionality, enabling AI/ML developers to use existing or new MLflow-enabled models to train within an EscrowAI TEE. The training output is delivered to an embedded, project-specific MLflow tracking server. Results are presented by algorithm version (experiment) and runs. For instance, a model v1 can be trained and tuned using hyperparameter permutations, and then a separate model v2 can be trained and tuned. Data from these runs can be tracked and compared within the embedded MLflow instance to select the optimum model-hyperparameter mix. Trained models are stored within EscrowAI and are available for download.

Validation, inference, query

EscrowAI includes important AI/ML life cycle capabilities. The platform allows for secure algorithm operation on protected data within a TEE. EscrowAI checks the algorithm output for PHI/PII before it leaves the TEE using a customized policy that the collaborating parties approve before the run initiation. The algorithm is not restricted, but the output must match the policy. The algorithm can

run data queries, statistical analysis, produce inferences or produce algorithm performance metrics.

These tables list the full functionality of EscrowAI 2.0.1, which is available through one or more license options. EscrowAI project functionality is activated by a Collaboration Unit (CU), a usage license that enables the project collaboration functionality for a fixed duration. Either party, or both, can contribute CUs to a project. Depending on duration, a project may need one or more CUs.

Itemized features[1]

Project AI lifecycle
Model training • Integrated MLflow training environment with project-specific tracking server enabling full performance monitoring by algorithm version and hyper-parameter run. • MLflow standardized and custom reporting • MLflow graphical output • Model repository and download • Algorithm test, validation, and query • Core functionality
Algorithm test, validation and query • Test and validate algorithms with custom performance reporting • Query data sets (statistical, quality checks, demographic summaries, etc.) • Core functionality

Core Functionality

Workflow
Organization Homepage with role-based Project listing.
Project collaboration workspaces.
Chat and notification within Projects.

[1] Available features may vary based on license type.

Choice of Trusted Execution Environment, by Project (subject to Azure Region availability) • Confidential VM (AMD SEV-SNP based) in Azure • VM-isolated Confidential Containers on Azure Container Instances • Confidential Virtual Machines with Confidential GPU
Configurable workloads.
Multiple-container compute capability within the Trusted Execution Environment using Docker Compose.
Algorithm containerization and build for TEE execution.
Use GitHub actions to push algorithms to EscrowAI
Pull an algorithm container directly from the customer's Container Registry
Encryption utility (local) for key generation, key wrapping, and data/algorithm encryption (AES256)
Push-button initiation of the workload in the TEE.
Automatic TEE decommissioning.
Within project communications.
Inference / Analytics standardized reporting.
Report content checking and enforcement policies.
Immutable version and transaction recording in Azure Confidential Ledger
Artifact preservation.
Run logging, monitoring, and error messaging.
Customizable workload (application) logging
Estimated cloud compute cost tracking
Developer tools
EscrowAI Sandbox for testing before deployment into a TEE, supporting both validation and training cases
GitHub Integration - Support to upload algorithms using GitHub Actions
Automated Docker container build
Zero-trust & Security
Identity and access management: organization and role-based access
Single Sign On integration

MFA required for all users
All workloads process in Trusted Execution Environments
Attestation of Trusted Computing Base
PCR validation of compute elements
End-to-end encryption for algorithm and data, including in process.
Automated creation of output checking policies (Preview)
Memory isolation
FIPS 140-2 level 3 HSM Key Manager (third party managed)
Policy-based output checking
Third-party application security testing
Customer Support
Onboarding service
Online Help Center 24/7
Online Issue reporting and tracking 24/7
Customer Service Response 8/5 (PT)

Licenses

Name	***Description***
Collaboration Unit - 1:1 Validation	Collaboration Unit for a 1:1 validation project (includes test and query). This license supports a project between one data controller and one algorithm developer organization.
Collaboration Unit - 1:1 Training with MLFlow	Collaboration Unit for a 1:1 training project (using integrated MLflow, including 1:1 validation). This license supports a project between one data controller and one algorithm developer organization.
EscrowAI cloud tenant license - annual	Onboard a data controller cloud tenant for integration with EscrowAI DevOps automation, including annual support of changes to cloud tenant configuration.

BeeKeeperAI offers collaboration unit licenses through our academic and not-for-profit research support program. License bundles for use in research challenges are also available. Contact Sales for information.