SciDataCon 2025

Contribution ID: **283**

Type: **Poster**

Privacy-Enhancing AI-based Whole Slide Image Analysis

Monday 13 October 2025 19:10 (20 minutes)

Whole-slide images (WSIs) drive state-of-the-art computational pathology, but hospitals typically restrict their analysis to isolated, air-gapped workstations because these gigapixel slides contain highly sensitive patient data. On such systems the workflow for a single case is onerous: (i) technicians copy the multi-gigabyte WSI to a removable medium and walk it to the secure workstation; (ii) the slide is partitioned into patches (≈7 min); and (iii) deep-learning inference —runs for ≈20 min. With sequential processing and manual hand-offs, throughput stalls well below the 50 cases per day target for routine diagnostics.

We present a privacy-preserving, cloud-enabled pipeline that removes the physical-transfer bottleneck while maintaining strict confidentiality guarantees. The solution hinges on hardware-based trusted execution environments (TEEs):

- TEE Encryptor (on-premises). After verifying a remote enclave's attestation, it establishes an ephemeral session key, preprocesses each WSI locally, slices it into patches, encrypts the patches with the session key, and transmits them over TLS.
- TEE Analyzer (cloud). Hosted in an AMD SEV or Intel TDX enclave, it decrypts patches only inside protected memory, executes the two-stage deep-learning cascade, re-encrypts results with the same session key, and stores all artefacts in an object store accessible solely within the private analysis service. Encrypted results return to the hospital and are decrypted by the TEE Encryptor.

Because computation now runs on elastic cloud hardware, multiple TEE Analyzer instances can be launched in parallel. A deployment with ten enclaves cuts effective turnaround to minutes per case and comfortably exceeds the 50-case-per-day target, all without exposing WSIs or predictions in plaintext to the cloud operator. Regulatory mandates such as HIPAA and GDPR are satisfied because the data never leave trusted memory or the hospital in unencrypted form.

Our results show that confidential-computing clouds can deliver order-of-magnitude improvements in digital-pathology throughput while preserving patient privacy. The proposed architecture decouples sensitive data from untrusted infrastructure, offering clinicians a scalable, secure path to real-time, image-centric diagnostics.

Primary authors: Mr TAN, Benjamin Hong Meng (A*STAR - Institute for Infocomm Research); Dr CHENG, Chee Long (SingHealth); Mr CHAO, Jin (A*STAR - Institute for Infocomm Research); Ms AUNG, Khin Mi Mi (A*STAR - Institute for Infocomm Research); MIGUEL, Rodel (A*STAR - Institute for Infocomm Research)

Presenter: MIGUEL, Rodel (A*STAR - Institute for Infocomm Research)

Session Classification: Poster Session

Track Classification: SciDataCon2025 Specific Themes: Infrastructures to Support Data-Intensive Research - Local to Global