



Introducing 221b

On-Premise Agentic AI & Data Intelligence Company


Introducing our company, our technology, and what we can do for regulated industries.



Controlled AI, Verified Data — A
Governed Workflow for Regulated
Industries

About Us: A Hong Kong-Born AI Company Adopted by Government

From Cyberport inception to live deployment in law enforcement — five years of progress



ADLER CURRENT STATUS — 5 ACTIVE PROOFS-OF-CONCEPT IN HONG KONG

1 Law Enforcement Agency · 2 Banks · 1 Law Firm · 1 Critical Infrastructure Operator

Recognised by: Cyberport · HK RegTech Assoc. · IFTA

Selected to Represent Hong Kong AI at the Global Stage

World AI Conference (WAIC) 2026, Shanghai · 17-20 July 2026 · Hong Kong Pavilion

One of 18 Companies Representing Hong Kong



Selected through InvestHK and three Hong Kong science parks under the Finance, Legal & Cybersecurity category. Cyberport brought only six AI companies — 221b was one of them.

World's Largest AI Summit

WAIC 2026 attracted over 1,100 enterprises and 3,000+ innovations. 221b demonstrated Agentic AI and data intelligence live to global investors and industry leaders.

Aligned with National AI Policy


Prof. Sun Dong (GBS, JP), Secretary for Innovation, Technology and Industry, visited the Hong Kong Pavilion and met with participating companies.



Third-party validation: 221b's AI was chosen by Hong Kong's innovation ecosystem to represent the city on the international stage.

Observation, Adler and Radix

Three engines from the 221b product suite — the ones closest to daily regulated workflows



What is an Ontology?

A structure that tells the system "what this is" and "how things are related" — the same data becomes different once meaning is added

✗ Without Ontology

Data exists in separate systems — no connections between them

Customer Data Table			
ID	Name	Address	Phone

✗ — No connection between tables

Transaction Data Table			
Txn ID	Amount	Date	Account ID

✗ — No connection between tables

Alert Data Table			
Alert ID	Rule Code	Score	Account ID

✗ — No connection between tables

Suspicious Transaction Report Table			
Report ID	Status	Filed Date	Customer ID


The analyst sees only:

Account 12345 has one USD 50,000 transaction. Alert triggered by Rule 7.2. Score: 85. Nothing else — no background, no relationships.

Analysts must manually search every system to find connections.

✓ With Ontology

Same data — now meaningfully interconnected



The system sees:

Account 12345 (Zhang Dawen) transferred USD 50,000 to Account 67890, which shares an address with a flagged high-risk entity. Zhang Dawen is a director of HongYu Company, the registered owner of Account 67890. Device X also logged into Account 99999, which is linked to STR-2023-015.

System automatically identifies four layers of risk relationships.

System follows relationships to show what is connected and why.

Ontology = a definition map: what entities exist, what attributes they have, and how they are connected.

Same data. Same systems. Ontology adds meaning and relationships — the system now knows what else to investigate.

Built for High-Stakes, Regulated Workflows

Adler is deployed across five core use cases where speed, accuracy, and auditability are non-negotiable



Anti-Money Laundering

Automate transaction monitoring, entity resolution, and STR preparation — cutting analyst hours from days to hours with full audit trails.

Observation + Adler



Anti-Fraud

Cross-reference behavioural signals, device data, and relationship networks in real time to surface hidden fraud rings.

Observation + Adler



Procurement Compliance

Screen vendors, beneficial owners, and related parties against sanctions, litigation, and adverse media before contract award.

Adler + Radix



Legal Research & Due Diligence

Automate counterparty, director, and shareholder background checks using verified Mainland China court, asset, and tax records.

Adler + Radix



Physical Security

Aggregate and correlate access logs, incident reports, and personnel data to detect anomalies and support investigations.

Adler

All use cases run on a governed, auditable pipeline — every step logged, every output traceable.



Deployment: On-Premise, Controllable, and Compliance-Ready

Open-source, containerised architecture — deployable in your own data centre, private cloud, or air-gapped environment

Deployment Mode	On-Premise Bare-metal servers, private virtualisation, air-gapped environments	Cloud (Public / Private) AWS, Azure, Alibaba Cloud, GCP	Containerised Docker containers, Kubernetes orchestration
Model Platform	AWS Bedrock Managed foundation model service	Alibaba Qianwen Enterprise-grade large model service	Private Model Engine Self-hosted models, fully private data pipeline
Hardware Acceleration	NVIDIA GPU CUDA-optimised		Huawei Ascend NPU CANN architecture
No Vendor Lock-in Fully Containerised Open-Source Architecture No Additional Licensing			

Radix: Exclusive China Data Products

Four data products, two group-exclusive — queried by precise ID number, not probabilistic name matching

National Fixed Asset Data

GROUP EXCLUSIVE

The only market service offering query by unique identifier for every fixed asset — including the registration and preservation number for asset freezing and seizure — covering 2,500 fixed asset registration centres across Mainland China.

National Litigation Data

Precise query using the subject's ID number, replacing the market norm of "fuzzy name matching" — significantly reducing errors from same-name cases, with no reliance on self-reported information.

Relationship Network Background Check

Maps connections between individuals and companies — shareholding, directorship, guarantorship — verified by ID number, presenting a reliable map of a subject's actual commercial footprint in Mainland China.

Corporate Tax Data

GROUP EXCLUSIVE

The only market service with access to historical tax filing data — enabling estimation of actual profitability, cross-referencing financial statements against tax paid, and identifying undeclared offshore income.

Precise ID-based query + exclusive registry data = verifiable answers, not probabilistic name matching.

Thank You

We welcome your questions, thoughts, and feedback on the above.



Controlled AI, Verified Data — A
Governed Workflow for Regulated
Industries