

NUS ASIAN INSTITUTE OF DIGITAL FINANCE

The Grammar of Finance

Why Ontology Graphs Are the Missing Layer Between Financial Data and Reliable AI Agents

Dr. Joey Tan

Associate Professor, NUS Asian Institute of Digital Finance

October 2026

Abstract

Banks that put AI models to work on their own data often meet the same problem. The models draft well but fail on the questions that involve money: which clients are exposed to a counterparty, whether two records describe the same company, whether a figure means the same thing in two systems. A common cause is that the bank has never written down what its data means in a form software can follow. Such a description is called an ontology, and without one a bank cannot safely let AI agents act on its behalf. Singapore's public record shows the gap with no AI involved: the Hin Leong collapse, in which one cargo supported three financings from six banks; MAS's 2025 finding, after the S\$3 billion money-laundering case, that most of nine penalised institutions had anti-money-laundering policies but applied them poorly or inconsistently; and the shared definitions Singapore's banks have agreed with each other since 2020 to close gaps of this kind. Recent studies find that newer AI models answer more database questions correctly but still give confident wrong answers where definitions are missing, and that AI agents working from a fictional bank's policy documents completed about a quarter of their tasks. No published study has yet tested an ontology inside a bank. Most of the cost of building one lies in agreeing definitions, and the best public costing of related work suggests that the return comes from reuse rather than from the first project. Ten questions let a board test its own position, starting with one that costs nothing. One development would undo the case: AI able to work out a bank's ontology from its existing systems.

Keywords

Ontology graphs · Knowledge graphs · GraphRAG · Agentic AI · Data governance · Entity resolution · Trade finance · Anti-money laundering · Risk data aggregation · Singapore · Banking

Contents

Abstract	2
Executive Summary	4
1. The Cargo That Was Financed Three Times	5
2. What the Gap Already Costs a Bank	7
3. What Better AI Models Have and Have Not Fixed	9
4. What an Ontology Is, in Banking Terms.....	12
5. What Singapore Banks Have Already Built	14
6. Ten Questions for the Board	17
7. What This Means by 2030.....	20
Appendix: The Terms in Plain Language	22
References	24
About	30

Executive Summary

Banks that have put AI models to work on their own data often hit the same problem. The models draft text well but get wrong the questions that involve money: which clients are exposed to a counterparty once ownership is traced to the ultimate parent, whether two records describe the same company, and whether a figure in one system means the same as its counterpart in another. Errors of this kind make the answers hard to sign off. The usual response is to wait for a better model. More often the cause lies in the bank, which has never described its own business in a form a machine can use.

Such a description is called an ontology. It sets out what a bank deals in (clients, groups, facilities, collateral, limits) and the relationships and rules that connect them, precisely enough for software to follow. Filled with the bank's own records, it becomes a knowledge graph. If the ontology is the grammar of a bank's data, the knowledge graph is the text written in it.

IN PLAIN TERMS

Three terms

- **An ontology** is a written description of the things a bank deals in and how they relate, precise enough for a computer to act on. It works like a dictionary that also states the rules.
- **A knowledge graph** is that description filled with the bank's actual records. The links between a client, its parent, its facilities and its collateral are stored as facts instead of being worked out again for each question.
- **An AI agent** is software that acts on a request instead of only answering it. It plans a few steps, uses other systems on the user's behalf and ends by changing something, such as releasing a payment, closing an alert or updating a record.

The problem is narrower than poor data quality. A bank's data can be accurate and complete and still be unusable by an AI system, because nothing records what each item means or how one record relates to another. Until that is written down, a bank cannot safely let software act for it outside a demonstration. Many controls a board relies on, such as limits on exposure to a corporate group, already assume this work has been done. Supervisors' assessments of banks' risk data suggest that in many banks it has not.¹

Two Singapore cases, neither involving AI, show what it costs when rules exist on paper but systems cannot apply them consistently. In 2020 a single cargo of diesel supported three separate financings from six banks, each holding documents that appeared to be in order.² In July 2025, after examinations connected with the S\$3 billion money-laundering case uncovered in 2023, MAS imposed composition penalties of S\$27.45 million on nine financial institutions. Most of them had established anti-money-laundering policies and controls, and MAS found that the breaches arose from poor or inconsistent

¹ Basel Committee on Banking Supervision, November 2023.

² Global Trade Review, 19 August 2020, summarising the interim judicial managers' report.

implementation.³⁴ Where the industry has tackled problems of this kind, it has agreed shared definitions: a registry that tells a bank when another bank is financing the same document, and a platform on which six banks exchange financial crime information on terms MAS sets.⁵⁶

None of these responses is called an ontology. They matter because an AI agent depends on the same thing: rules and definitions that a system can apply. An agent also acts faster than staff, often in a process from which the person who would have caught a mistake has been removed.

Banks should now do internally what Singapore's banks have been doing with each other since 2020: agree what things are, write the definitions down and keep them current. Supervision points the same way. MAS has said that its forthcoming guidelines on AI risk management will cover AI agents.⁷ In July 2026 it published a voluntary industry framework under which no action by an AI agent is carried out until it has been declared, authorised and assessed.⁸ A board can begin at its next meeting, at no cost, by asking to see how one reported number was produced.

THE EVIDENCE SO FAR

What the studies show

- **About a quarter of tasks.** In a March 2026 study by Sierra, a maker of customer-service agents, and Princeton University, AI agents serving the customers of a fictional bank from about 700 linked policy documents completed about a quarter of their tasks correctly. They failed mainly at finding the right policy and applying it.
- **From 31 to 46 per cent.** GPT-4's accuracy on an academic benchmark of 95 databases, one of them from a bank, rose by this much when it was given written definitions of the business terms.
- **An error instead of a wrong number.** When dbt Labs, which sells a layer of defined business terms, re-ran a 2023 insurance test in 2026, current models answered most questions correctly whether or not they used the layer. The difference lay in how they failed: the defined layer returned an error when it could not answer, while direct queries returned a plausible wrong figure.

Newer models have narrowed the accuracy gap. They still give confident wrong answers where definitions are missing, and agents still struggle to find and apply a bank's own rules. No published study has yet tested an ontology on a bank's own systems.

³ Monetary Authority of Singapore, 4 July 2025; AsiaOne, July 2025, on the link to the 2023 case.

⁴ Monetary Authority of Singapore, 4 July 2025.

⁵ Association of Banks in Singapore, 23 June 2023.

⁶ Monetary Authority of Singapore, 1 April 2024.

⁷ Monetary Authority of Singapore, 13 November 2025; written parliamentary reply, 5 August 2026.

⁸ Monetary Authority of Singapore, 3 July 2026 (SAFR).


1. The Cargo That Was Financed Three Times

Six banks financed one cargo, and none could see the others' files.

In March 2020 a cargo of 780,000 barrels of ultra-low sulphur diesel sat aboard a single vessel. It supported three separate financings: a Crédit Agricole letter of credit for a sale to Unipet; Société Générale financing, on letters of credit issued by Rabobank and DBS, for a sale-and-repurchase with Glencore; and Natixis and OCBC letters of credit for an arrangement involving Trafigura and Winson. The bills of lading behind all three named the same vessel and cargo, and differed only in date.⁹

MARCH 2020 · ONE VESSEL, ONE CARGO

THREE FINANCINGS, SIX BANKS



Each bank's documents appeared in order. None knew that the cargo behind its facility also secured two other financings. Reuters reported debts of about US\$3.85 billion to 23 banks. Over six years, trade finance failures of this kind left creditors facing about US\$9.26 billion in potential losses.

Figure 1. One cargo, three financings. Each bank's documents appeared in order, and nothing connected the three.

Each bank financed against documents that appeared to be in order. None could see that the cargo behind its facility also secured two other financings. No shared record linked a cargo to every financing against it, and each bank could see only its own books.

The interim judicial managers' report, filed in June 2020, found forged bank statements, bills of lading, sales contracts and invoices. It found inventory overstated by at least US\$809 million in the accounts for the year to October 2019, and export contracts, presented to fifteen banks in support of 58 import letters of credit, with no sale behind them.¹⁰ Reuters reported debts of about US\$3.85 billion to 23 banks; later reporting put the exposure at about US\$3.5 billion.¹¹ The founder, Lim Oon Kuin, was sentenced in 2024 to seventeen and a half years for cheating HSBC and instigating forgery. On appeal in March 2026 the term was cut to thirteen and a half years.¹² Over six years, commodity trade finance failures including Hin Leong, Agritrade, Hontop and Zenrock left banks and other creditors facing about US\$9.26 billion in potential losses. Bloomberg described the cases as involving alleged fraud, forgery, multiple pledges of the same cargo or questionable accounting.¹³

⁹ Global Trade Review, 19 August 2020.

¹⁰ Global Trade Review, 19 August 2020, citing the interim judicial managers' report of June 2020.

¹¹ Reuters, 16 April 2020, citing industry sources. The later figure is from Bloomberg, 9 July 2020.

¹² Singapore Police Force, 18 November 2024, on the sentencing; AsiaOne, 4 March 2026, on the appeal.

¹³ Bloomberg, 9 July 2020.

Forgery caused the losses. With no shared record of what each cargo already secured, the same barrels could be pledged again.

Treating this only as fraud and conduct risk misses a wider lesson. A control states a rule about the world, such as: do not lend twice against the same collateral. To apply it, a system needs collateral, a pledge and “the same” defined clearly enough to check, across records that each institution keeps differently. Every bank had the rule, and none could apply it beyond its own books. A control inside a single bank has the same weakness wherever two of its systems record the same thing in different ways.

No shared definition would have stopped forged documents being presented, and a determined fraudster can beat good systems. A shared record changes whether the fraud survives the next step, when the same barrels are offered to other banks. Double financing depends on the gaps between institutions.

An AI agent working inside a bank meets the same problem. When two systems record the same client differently, an experienced member of staff may notice. An agent working at speed, in a process from which that person has been removed, is less likely to. Questions about concentration, connections, exposure and identity all depend on how records join up. Answering them reliably requires a written statement of what each record refers to and how records connect, which is what an ontology provides.

2. What the Gap Already Costs a Bank

Five places where a bank already pays for data whose meaning was never written down, none of them labelled that way in the budget.

Hin Leong was an extreme case, and a board could argue that extreme cases are what capital is for. The same gap also costs money every year, in the five areas shown in Table 1.

Where the cost falls	How the gap causes it	Who carries it
Financial crime remediation	Risk ratings, source-of-wealth checks and transaction reviews depend on terms defined in policy documents but not in any system that applies them consistently	Compliance staff, penalties, and remediation required by the supervisor
Client and counterparty identity	The same company is recorded differently across businesses, systems and countries, so group exposure is assembled by hand and is out of date when it arrives	Credit risk, through understated group exposures
Regulatory reporting and reconciliation	Each regulator defines its fields separately, and few banks record which requests ask for the same thing, so each return needs its own extraction	Finance and reporting teams, every reporting cycle
Risk data aggregation	BCBS 239 asks for accurate, complete and adaptable risk data across systems that were never described in one place	A decade of spending on a standard that few of the largest banks fully met by 2023
AI that stops at the pilot	Demonstrations work and production stalls, often because letting software act for the bank requires a description of the bank that nobody has written	Written-off spending, and the credibility of the next business case

Table 1. Five costs with a common cause. No single project will recover all of them, and a board should test each line before accepting it.

The first line has recent public figures, and they trace back to one case. In August 2023 Singapore police arrested ten foreign nationals in a money-laundering investigation that came to involve assets of close to S\$3 billion. All ten were convicted.¹⁴ MAS examined the financial institutions that had dealt with the people involved and, in July 2025, imposed composition penalties totalling S\$27.45 million on nine of them, taking into account each institution's exposure.¹⁵ All nine had failed to corroborate significant parts of customers' source of wealth in higher-risk situations. Eight had not properly reviewed unusually large or inconsistent transactions. Five had customer risk assessment methods that were not robust enough to classify clients correctly.¹⁶ MAS observed that most of the institutions had established anti-money-laundering policies and controls, and that the breaches arose from poor or inconsistent implementation.¹⁷

Two bank employees were also prosecuted. In October 2025 a former relationship manager at Julius Baer was jailed for four months for using a forged document: he had guided a client in forging a Chinese tax payment certificate and submitted it to the bank to open her account. A former Citibank relationship manager was jailed for two years for forgery, money laundering and obstructing justice.¹⁸ In March 2026

¹⁴ Singapore Police Force, 18 November 2024.

¹⁵ AsiaOne, July 2025; The Star, 5 July 2025; Monetary Authority of Singapore, 4 July 2025.

¹⁶ Monetary Authority of Singapore, 4 July 2025.

¹⁷ Monetary Authority of Singapore, 4 July 2025.

¹⁸ Singapore Police Force, 24 October 2025; Malay Mail, 23 October 2025.

MAS barred the two from the financial industry for seven and sixteen years.¹⁹ In the first case, the account-opening check rested on a document and on the employee who vouched for it.

BOARD QUESTION

How far apart are your policies and your systems?

A policy is a statement of intent. To apply it consistently at volume, terms such as higher-risk situation, source of wealth and unusually large must be defined precisely enough for a system to apply them to a client's actual records. The system must then apply them the same way each time, in a way that can be shown afterwards. Ask where those definitions are kept. If the answer is a policy document and the judgement of experienced staff, the bank is relying on people to apply the policy consistently, which is where MAS found the nine institutions fell short.

Singapore has also tightened its rules on company ownership. Since 16 June 2025, companies must confirm the details of their registrable controllers every year, the definition of a nominee shareholder has been widened, and nominee directors and shareholders must be filed with ACRA, the company registry.²⁰ Only public agencies can see the full register, so banks still collect ownership information from their clients and check it themselves.²¹ The new rules make that information more likely to be current. A bank gains most if its own systems record ownership as dated facts linked to the client, rather than as text in an onboarding file.

The fourth line has a long record. BCBS 239, the Basel Committee's standard on risk data, was published in 2013, and the largest global banks were expected to comply from 2016.²² In November 2023 the Committee reported that only two of the thirty-one banks it assessed fully complied with all the principles, and pointed to fragmented IT, legacy systems and manual processes.²³ Seven years after the deadline, the best-resourced banks had not solved the problem. A team that promises this bank will do better should be asked what will be different. The most convincing answers are a narrower scope and a named owner.

Figures on failed AI investment come from outside banking. A 2025 report from MIT's NANDA initiative found that 95 per cent of organisations were getting no return from generative AI, and blamed tools that do not learn or adapt to how firms work rather than the quality of the models.²⁴ Gartner expects more than 40 per cent of AI agent projects to be cancelled by the end of 2027.²⁵ Both figures need discounting: the first rests on interviews with 52 organisations and a short survey, and has been disputed; the second is a forecast. A more useful number is the bank's own: the staff-days spent last year working out what data fields mean. Most banks could estimate it. Few boards have seen it.

¹⁹ Monetary Authority of Singapore, 17 March 2026.

²⁰ Companies and Limited Liability Partnerships (Miscellaneous Amendments) Act 2024.

²¹ Rajah & Tann Asia (2025), on the 2024 amendments and access to ACRA's registers.

²² Basel Committee on Banking Supervision (2013).

²³ Basel Committee on Banking Supervision, November 2023.

²⁴ MIT Project NANDA (2025).

²⁵ Gartner, 25 June 2025.

3. What Better AI Models Have and Have Not Fixed

Current models answer database questions far better than they did in 2023. They still give confident wrong answers where definitions are missing, and AI agents still struggle with a bank's own policies.

Many banks began with retrieval-augmented generation. The bank's documents are split into passages; when someone asks a question, the system finds the passages that most resemble it and passes them to the AI model to answer from. The method is cheap and quick to set up, and works well for tasks such as summarising a credit file, drafting a client note or finding the clause in a policy that covers a case.

It fails when the passages that most resemble a question are not the ones that answer it. Three kinds of question cause this, and all three are routine in banking.

- **Questions whose answer is in no single document.** Which clients are exposed to this counterparty, through which facilities and guarantees, once ownership is traced to the ultimate parent? The answer has to be assembled by following a chain of links.
- **Questions about the whole book.** Where is concentration building across businesses that each look comfortable on their own? Which patterns recur across this quarter's alerts? Retrieving the ten most relevant passages cannot answer a question about all of them.
- **Questions of identity.** One system holds an internal client code, another a Legal Entity Identifier (the global reference code for legal entities), a third a name typed by hand. Are they one company or three? The search returns all three and leaves the AI model to guess. A confident wrong guess looks the same as a right one, so nothing in the answer warns the reader.

An early measurement came in 2023 from researchers at data.world, which sold a data catalogue built on a knowledge graph and was bought by ServiceNow in 2025. On a test database built to an industry-standard insurance design, GPT-4, then a leading model, answered 16.7 per cent of business questions correctly when it wrote database queries directly, and 54.2 per cent when the same data was presented as a knowledge graph. In a 2024 follow-up, with each query checked against the ontology's rules and repaired where needed, 72.6 per cent of answers were correct and a further 8 per cent came back as unknown rather than wrong.^{26,27} Models have improved since, and these figures are out of date.

The test has since been re-run with newer models. In April 2026 dbt Labs, which sells a layer of defined business terms, reported that current models querying the original database directly answered 64.5 per cent of its 11 questions correctly, against 32.7 per cent in its own 2023 run. Through dbt's defined layer, accuracy rose from 60.5 to 72.7 per cent. After three additional data models were built, direct queries reached 84 to 90 per cent and the defined layer 98 to 100 per cent.²⁸ An August 2026 preprint by an author

²⁶ Sequeda, Allemang and Jacob (2023).

²⁷ Allemang and Sequeda (2024).


²⁸ dbt Labs, 7 April 2026.

with no stated link to either company used dbt's version of the test with 38 questions. A governed approach answered 37 of them correctly in all three runs; direct queries managed 21.²⁹

For a bank, how each approach fails matters more. dbt reports that its defined layer returned an error when it could not answer, while direct queries returned a plausible wrong number. The preprint counted the difference: across 114 runs, direct queries returned 29 wrong answers with no sign of error, while the governed approach returned none and declined to answer once.³⁰ An error message can be caught and escalated. A plausible wrong figure can reach a credit file, a regulatory return or a client before anyone notices.

THE SAME INSURANCE TEST, RE-RUN

2023 AND 2026



Share of questions answered correctly on the same insurance test (11 questions, each run 20 times), as run by dbt Labs, which sells the defined-layer approach. Where two current models differed, the range is shown. When the defined layer could not answer, it returned an error; direct queries returned a plausible wrong figure.

Figure 2. The 2023 insurance test, re-run by dbt Labs in 2026. Newer models closed much of the gap. The defined layer still scored higher and, when it could not answer, returned an error.

Academic studies point the same way. In BIRD, a benchmark of 95 databases across 37 fields, one of them from a Czech bank, giving GPT-4 written definitions of the business terms raised its accuracy from 30.9 to 46.4 per cent.³¹ Spider 2.0, built from real company databases that often have more than a thousand columns, showed how much harder such data is. One of the strongest models of late 2024 solved 21.3 per cent of its tasks, against 91.2 per cent on an older benchmark of simpler databases.³² A 2026 test by MotherDuck, a database company, looked at simple data. On eleven small databases averaging seven tables each, written descriptions added 1.1 percentage points on one set of questions and 0.2 on another, but 8.7 points on the hardest database.³³ Definitions add little where data is simple and most where it is

²⁹ Yi Ai (2026), preprint.

³⁰ Yi Ai (2026).

³¹ Li and colleagues (2023). On errors in the bank database, Wretblad and colleagues (2024).

³² Lei and colleagues (2025).

³³ MotherDuck, 6 February 2026.

hard to interpret. A bank's data, which supervisors describe as fragmented and partly manual, is likely to fall in the second group.

HOW FAR TO TRUST THESE NUMBERS

Much of this evidence comes from companies that sell the approach being tested: data.world, dbt Labs and MotherDuck. The re-runs used one insurance database with 11 to 38 questions; the preprint used dbt's version of the test and has not been peer-reviewed. BIRD and Spider 2.0 are academic, but neither tests an ontology as such, and a later study found errors in about half of the questions in BIRD's bank database. No published study has tested an ontology against a bank's own systems. The evidence points one way; how large the effect would be inside a particular bank is unknown.

These are failures of question answering. Banks are now being offered AI agents, which plan a series of steps, use other systems and change records. MAS has said that its forthcoming guidelines on AI risk management will cover AI agents as well as generative AI.³⁴

A recent test of AI agents in a banking setting comes from Sierra and Princeton University. In a study published in March 2026, AI agents served the customers of a fictional bank. They worked from about 700 linked documents describing products and policies, used banking tools to update accounts, file disputes and process applications, and were graded on whether the bank's records ended up correct. The strongest set-up completed about a quarter of tasks, and about 13 per cent when the same task had to be done correctly four times in a row. The authors traced the failures to two causes: finding the right document in a densely linked knowledge base, and reasoning correctly over complex internal policy.³⁵ On a public leaderboard for the same banking tasks, the leading models in October 2026 completed just over half.³⁶ Neither source tests whether an ontology would close the gap. Both show that the gap exists for software that acts in a banking setting.

Each of the three failures is more serious when an agent is involved. An assistant that confuses two clients gives a wrong answer, and the relationship manager reading it may notice. An agent that confuses two clients takes a wrong action: it releases a payment, clears an alert, approves a higher limit or files a report naming the wrong company. Business cases for agents often count the removal of that human check as the saving. A bank cannot count the saving and also rely on the check.

The agent software now on the market does not solve this. Orchestration tools manage the sequence of steps an agent takes and the hand-offs between agents; connection standards govern how an agent reaches another system. Neither defines what the bank's data or actions mean.³⁷ Buying them without that definition gives a bank faster ways to act on records it still cannot interpret reliably.

³⁴ Monetary Authority of Singapore, 13 November 2025.

³⁵ Shi and colleagues (2026).

³⁶ Artificial Analysis, accessed 1 October 2026.

³⁷ Orchestration tools include LangGraph, CrewAI and Microsoft Agent Framework; connection standards include the Model Context Protocol.

For a board, the studies support three conclusions. Newer models have narrowed the accuracy gap, so the case for written definitions cannot rest on accuracy figures alone. Where definitions are missing, AI systems still give wrong answers with confidence, and in the tests so far a defined layer turned many of those into errors that could be caught. Agents working from a bank's policy documents fail mainly because they cannot find and apply the right policy, which is the problem an ontology is designed to address. Whether it does so inside a real bank has not been tested.

4. What an Ontology Is, in Banking Terms

What an ontology contains, and how it differs from the classification schemes a bank already runs.

The standard definition, from 1993, calls an ontology “an explicit specification of a conceptualization”.³⁸ For a bank, it sets down four things: the kinds of thing that exist, such as client, facility, collateral and guarantee; the details recorded about each; the named relationships between them, such as is controlled by and is secured on; and the rules that follow, such as: if A controls B and B controls C, then A controls C. Loading the bank's records into that structure produces a knowledge graph.


Banks already run something that looks similar but does a different job: a taxonomy, or classification scheme. A taxonomy files things. A facility is classed as Corporate, then Secured, then Commodity trade finance, which lets the bank report how much exposure sits in each class. A taxonomy cannot show that the borrower is 62 per cent owned by a parent that has separately guaranteed two loans in another business line, so that the bank's exposure to the group is higher than the limit system shows. An ontology records that connection and what follows from it.

³⁸ Gruber (1993).


THE SPINE

A FEW DOZEN CONCEPTS, THREE RULES

The concepts a bank already uses



What gets said about how they connect



The three rules that carry most of the leverage



Anyone who has written a credit policy or set up a limit framework will recognise these terms. The work lies in agreeing on them across the bank and writing down what follows from them.

Figure 3. The core of a bank ontology: the things a bank already deals in, how they connect, and three rules.

The core is likely to be small: a few dozen concepts, such as clients, legal entities, groups, beneficial owners, facilities, instruments, collateral, guarantees, limits, policies and controls. Anyone who has written a credit policy will recognise the list. No bank appears to have published its enterprise ontology, so this is a judgement drawn from published industry models such as FIBO; there is no case a board can inspect. The practical test inside a bank is whether its own people can agree the list, and how long that takes.

Three rules do most of the work.

- 1. Control passes up the chain.** If A controls B and B controls C, then A controls C. Group exposure can then be calculated from the ownership chain, instead of being read from a field that someone updates by hand and may forget after a restructuring.
- 2. One client, one identifier.** The same company is recognised as the same company across businesses, systems and countries. The decision that two records refer to the same party is recorded as a fact, with a date and an author, and can be reviewed. Often it is instead a matching score from an overnight batch that is rarely reviewed.
- 3. Every definition has a version.** A reported figure can be traced to the definition in force on the day it was produced. Changing a definition becomes a dated, approved event instead of an unnoticed break in the numbers.

The third rule speaks most directly to supervisors. Policies are revised and risk ratings recalibrated, and figures produced under different definitions cannot safely be compared over time. Versioned definitions let a bank manage those changes the way it manages new releases of its software, and they answer a question supervisors ask after an incident: what was the bank's rule on the date in question?

Four practical benefits follow.

- **Identity decisions can be reviewed.** The decision that two records are the same client carries a date and an author, so it can be challenged, corrected and audited.
- **Figures come from the records.** The AI model writes the query and explains the result, but the number is calculated from the bank's own records, and the query can be inspected.
- **The explanation is built in.** The route from a facility, through the borrower, to the group, under a given policy and version, is itself the audit trail.
- **Access rules can follow the data.** Permissions can be attached to the records themselves, so an agent sees only what the person it acts for is allowed to see. Without that control, an agent should not be allowed near client data.

A bank does not have to start from nothing. The Financial Industry Business Ontology (FIBO) is a free, published description of financial concepts, updated quarterly by ontologists and industry specialists. The EDM Council, which maintains it, lists consistent regulatory reporting, including under BCBS 239, among its uses.³⁹ It is a starting vocabulary rather than a finished model. No industry model knows what a particular bank means by a relationship, a limit or a connected party, and agreeing that is most of the work. Starting from FIBO saves the bank writing the common parts itself.

5. What Singapore Banks Have Already Built


Four efforts that are not ontologies, and what they show.

Singapore's banks have done a version of this work several times, each time because no single bank could answer a question alone. None of it produced an ontology. Public accounts suggest that in each case the slow part was reaching agreement between institutions, on definitions and on the terms of sharing, rather than building the software.

³⁹ EDM Council, FIBO.

CROSS-BANK ARRANGEMENTS IN SINGAPORE

FOUR STEPS SINCE 2020



None of these is an ontology. The first three rest on agreements about what something is (a financed document, a warning sign), reached because no single bank could answer the question alone.

Figure 4. Cross-bank arrangements in Singapore since 2020. None was built as an ontology.

A registry that tells a bank another bank is financing the same document

After the 2020 frauds the industry acted in two stages. In October 2020 Standard Chartered and DBS led a trial that let a bank check whether another bank had already been asked to finance a given title document, by comparing encrypted data so that the underlying records never left each bank.⁴⁰ In June 2023 the Association of Banks in Singapore launched the Trade Finance Registry, with UOB as chairman bank. UOB, Standard Chartered, DBS, Citibank, OCBC and BNP Paribas led the project, and some thirty other banks joined. Banks register their corporate customers' transactions in encrypted form and are told in near real time if another participant registers the same one.⁴¹

In practice, the registry tells a bank that the document it is about to finance is already being financed elsewhere. It rests on one thing the participating banks agreed: what counts as the same transaction. More than thirty banks take part, and the registry addresses one of the gaps exploited in the 2020 failures. It needed no AI.

A platform where the regulator sets the warning signs

COSMIC, launched by MAS in April 2024, lets six banks (DBS, OCBC, UOB, Citibank, HSBC and Standard Chartered) share information about customers who show signs of financial crime in three areas of risk: misuse of legal persons, misuse of trade finance and proliferation financing.⁴² Sharing is allowed only when a customer shows objectively defined warning signs. When Parliament debated the enabling law in May 2023, the government said MAS would issue a directive setting out the threshold for each risk and the

⁴⁰ Trade Finance Global, 8 October 2020.

⁴¹ Association of Banks in Singapore, 23 June 2023.

⁴² Monetary Authority of Singapore, 1 April 2024.

warning signs that go with it.⁴³ The banks therefore share information on terms set by the regulator rather than on each bank's own view of what looks suspicious.

Under COSMIC, banks may share information about a customer only when warning signs defined in advance are present.

If six banks need definitions set by the regulator to share information with each other, each might ask what it uses internally, across private, corporate and retail businesses that may each have taken on the same client separately.

What the analytics achieved, and its limits

In 2017 OCBC said it had become the first Singapore bank to use artificial intelligence and machine learning against financial crime. Tested on a year of corporate banking transactions, its system reduced alerts that needed no further review by 35 per cent and made the identification of suspicious transactions more than four times as accurate.⁴⁴ The result was real, but it is now nearly ten years old, and it shows the limit of such tools: they worked within one bank's view of its own transactions.

Research since then points the same way. A review of 97 studies found that analysing the connections between accounts improves on conventional methods. It also warned that newer techniques such as graph neural networks need care on real laundering data, where the cases being sought are a tiny fraction of the whole, and that results on synthetic data can be too optimistic.⁴⁵ For a board, the lesson is that analysing connections helps, that published improvements should be discounted, and that results depend on the quality of the connections being analysed.

MAS's 2026 work shows both the direction of travel and the obstacle. In May 2026 it announced a trial with five banks, the Government Technology Agency and the Singapore Police Force to detect scams earlier using the banks' combined data, with account numbers scrambled so that only the bank holding an account can identify it.⁴⁶ The privacy protections are described in detail; how the five banks' data is made consistent is not. A September 2026 trade press analysis identified the main practical challenge as aligning the banks' transaction formats, customer risk indicators and fraud categories closely enough for models to make reliable comparisons. It also reported MAS's managing director, Chia Der Jiun, as saying that a model that performs well at one bank may behave differently at another.⁴⁷ Any bank joining a pooled-data exercise should ask who is doing that alignment, and whose definitions are being used.

⁴³ Monetary Authority of Singapore, 9 May 2023, second reading speech on the amendment bill.

⁴⁴ OCBC Bank, 7 November 2017.

⁴⁵ Deprez and colleagues (2024).

⁴⁶ Monetary Authority of Singapore, 4 May 2026.

⁴⁷ QA Financial, 22 September 2026.

One bank's own account

A DBS senior vice president has described a framework that uses a graph to capture “some of the knowledge that we have in our bankers' heads”. It combined application, transaction, payment, credit bureau and company registry data, and analysed the connections for product recommendations, supplier identification and concentrations of risk across directors, shareholders, industries and geographies.⁴⁸ The source is limited: one executive speaking at a software vendor's conference in 2022, as reported by the vendor. DBS has not published an enterprise ontology. The account shows only that a large Asian bank found the questions it wanted answered were about connections, and chose to store its data as connections to answer them.

WHAT THESE CASES DO NOT SHOW

None of these four is an ontology, and none proves that building one pays. The Trade Finance Registry settles a single relationship. COSMIC is an information-sharing platform governed by thresholds and a list of warning signs, closer to a controlled vocabulary than to a model with rules. The OCBC result comes from a 2017 trial in one bank. The DBS account is a conference talk. The public record supports a narrower conclusion: when the industry has needed to answer a question no single bank could answer, the slow part appears to have been agreeing what things are and how to identify them. A director looking for proof that an internal build will pay off will not find it in the public record, and should price that uncertainty into the decision.

6. Ten Questions for the Board

Five to ask inside the bank and five to ask anyone selling it an AI agent. The first of each costs nothing.

A board may reasonably ask why it should act when the evidence that an internal ontology pays is incomplete. There are three reasons. The first steps cost little, needing time from senior people rather than a budget. The supervisory direction is set. MAS has said its AI risk guidelines will apply to AI agents, and its July 2026 framework requires each action by an agent to be declared and authorised before it runs. That presupposes the bank can state what each agent may do.⁴⁹ And the bank already pays for inconsistency in how its terms are applied, in remediation and at worst in penalties, whether or not it ever deploys an agent.

It is a board matter because it concerns who is accountable for what the bank's terms mean, a question a technology committee cannot settle.

⁴⁸ Neo4j, 2 January 2023, reporting remarks made at GraphSummit Singapore 2022.

⁴⁹ Monetary Authority of Singapore, 3 July 2026 (SAFR).

INWARD**Five questions for your own bank**

- **Ask to see how a reported number was produced.** Pick a figure that went to a regulator or a credit committee last quarter. Ask for it to be reproduced with each step shown: the source system, each definition applied, the final number. How long that takes, and how difficult it proves, is likely to tell the board more than a maturity assessment.
- **Count your definitions of client.** How many working definitions of client, group, exposure and connected party exist across the businesses, the limit system, the regulatory returns and the financial crime policy? Often nobody is responsible for that number.
- **Give the definitions an owner and a release schedule.** A named owner, a version number, a change log, and a rule that changing a definition is a dated, approved event. Banks already control changes to their software this way and can apply the same discipline to what their data means.
- **Put the controls in the budget.** A written list of the actions each AI system may take; access that follows the user's own permissions; a check on each action before it runs; a standing set of test cases; and a record of what was done and why. How much of this a system needs depends on how much damage a wrong action could do. Anything that can move money needs all of it.
- **Ask what the second project will cost.** Any team or vendor can make a case for the first use. Ask what the second and third will cost, and why. If the answer is about the same as the first, nothing reusable is being built, and the bank is paying for a run of pilots rather than a capability.

One public study has costed reuse, though not inside a single bank. In a 2020 pilot of digital regulatory reporting, the Financial Conduct Authority and the Bank of England gathered cost data from eleven large lenders and 126 smaller firms. They found that “the lack of standardisation of the definitions and descriptions of data by firms is a significant barrier”.⁵⁰ They then modelled the payback. A single reporting area built on new central infrastructure would not pay back within ten years. The same area reusing existing industry standards would pay back in three years, and three areas introduced one after another in seven.⁵¹ They concluded that one area on its own was unlikely to be viable, that the approach warranted further investment, and that it should proceed in small steps.⁵²

The study has limits. It covers regulatory reporting across an industry, not an ontology inside one bank, and the pilot itself was not put into production. No bank appears to have published what building an ontology costs. The study's pattern, in which the return comes from reuse rather than from the first project, is a reasonable expectation; any precise figure a vendor offers should be treated as unsupported.

⁵⁰ Financial Conduct Authority and Bank of England (2020).

⁵¹ Financial Conduct Authority and Bank of England (2020).

⁵² Financial Conduct Authority and Bank of England (2020).

OUTWARD**Five questions for anyone selling you an AI agent**

- **What is it allowed to do, and can you list it?** If its possible actions cannot be written down, its risk is hard to assess, and MAS's proposed guidelines expect banks to keep an inventory of their AI systems and assess the risk of each.
- **Show me it getting something wrong.** Ask the vendor to show the system declining to answer. A system that never declines may not have been designed to recognise the limits of what it knows.
- **Whose definition of client is in the demonstration?** If the demonstration works because the vendor used its own data model, the board has been shown someone else's bank.
- **Can we take our definitions with us?** If the bank's definitions of client, exposure or connected party can only be read inside the vendor's platform, the bank depends on that vendor for its own vocabulary, and changing vendor will mean rebuilding it.
- **Who is accountable when it is wrong?** Expect a named person on the bank's side. MAS's proposed guidelines place accountability with the board and senior management, and the EU AI Act sets documentation, traceability and human-oversight duties for credit-scoring systems.

Supervision is making the inward questions harder to put off. MAS's proposed guidelines on AI risk management would require an accurate, current inventory of AI systems, an assessment of each one's risk based on its impact, its complexity and the bank's reliance on it, and controls in proportion, with the board and senior management accountable.⁵³ The risk of an agent whose permitted actions are written down can be assessed; the risk of one limited only by what a user might type into it is much harder to judge. The industry framework for AI agents that MAS published in July 2026 goes further. It places a check between every decision an agent makes and the action that follows, so that nothing is carried out until it has been declared, authorised and assessed.⁵⁴ The framework is voluntary and states that it is not supervisory guidance. Its checks are hard to apply unless a bank can say what each agent may do and what the data it acts on means, which is what an ontology records.

One way to sequence the work: in the first ninety days, run the reporting test, count the definitions, and choose one decision whose owner will take responsibility for its definitions, such as a limit check, an alert type or a step in client onboarding. Over the next two or three quarters, build the core ontology from the systems of record, add the checks to the process, and run one AI agent with a short, written list of permitted actions. In the second year, extend the work by reusing what was built. The first project should be judged by how much less the second one costs.

Most of the cost is people's time. The work needs a business owner senior enough to settle disputes about definitions; the bank's experts in credit, financial crime and finance, who know what the terms mean in

⁵³ Monetary Authority of Singapore, 13 November 2025.

⁵⁴ Monetary Authority of Singapore, 3 July 2026 (SAFR).

practice; and a few specialists who can write those meanings down in a form software can use. Specialists can be hired. The owner and the subject experts cannot.

7. What This Means by 2030

Four developments to plan for, and one that would undo the case for an ontology.

Forecasts about AI have a poor record, and none of these developments is certain.

Supervision moves from principles to checks

MAS has proposed a twelve-month transition once its AI risk guidelines are issued; they had not been finalised by October 2026.⁵⁵ Under the EU AI Act, obligations for high-risk systems, which include credit scoring, now apply from December 2027, after the EU postponed them in July 2026.⁵⁶ Singapore has already turned one set of principles into checks. The FEAT principles, developed with the industry in 2018, were followed by the Veritas programme, which has published assessment methods, case studies and open-source code for banking and insurance uses, including credit scoring and fraud detection.⁵⁷ Expectations that a system can check are easier for a bank to meet and for a supervisor to verify. A bank that can express a control as a question its systems can answer is likely to show compliance faster than one that can only describe the control in a document.

Banks will be asked to connect to more shared systems

The trade finance registry and the information-sharing platform both require a bank to supply data on definitions agreed with others, and the pooled-data trial will need the same.^{58,59} Payments have already been through this. The old message formats for cross-border payments were retired in November 2025 in favour of a structured standard, ISO 20022, more than twenty years after it was first published.⁶⁰ Each new connection is cheaper for a bank that has defined its own data than for one that has not, because the alternative is a separate translation for every connection, built and maintained by hand. Regulators elsewhere are moving the same way. In June 2026 the European Central Bank set out the timetable for its Integrated Reporting Framework, which will bring euro area banks' statistical reporting into a single standardised framework, with official reporting due to start in the second quarter of 2031.⁶¹ Alongside it, central banks and commercial banks maintain a free, voluntary dictionary in which the concepts behind reporting obligations are “identified and described once”.⁶² It is a shared definition of reporting data, built by supervisors and banks together.

⁵⁵ Monetary Authority of Singapore, 13 November 2025; written parliamentary reply, 5 August 2026.

⁵⁶ Regulation (EU) 2024/1689, Annex III; Regulation (EU) 2026/1744.

⁵⁷ Monetary Authority of Singapore, FEAT principles (2018) and Veritas Initiative.

⁵⁸ Association of Banks in Singapore, 23 June 2023.

⁵⁹ Monetary Authority of Singapore, 1 April 2024.

⁶⁰ Payment Expert, 21 November 2025.

⁶¹ European Central Bank, 8 June 2026.

⁶² European Central Bank, Banks' Integrated Reporting Dictionary.

Keeping definitions current will be the hard part

The Legal Entity Identifier shows the risk. At the end of 2025 there were more than 2.93 million active identifiers, but in the last quarter of the year only 56.7 per cent had been renewed, which means their details had been re-checked within the past year.⁶³ For more than two in five, the details may be out of date. A definition needs an owner, a version and a maintenance budget. A bank that pays to build its definitions but not to maintain them will end up with a set nobody trusts. An out-of-date definition can be worse than none, because people still rely on it.

Size may not be an advantage

This point rests on reasoning rather than evidence. If most of the work is reaching agreement inside the institution, a mid-sized bank with one credit policy and one client database may define its business faster than a large bank with many of each, because agreement gets harder as an organisation grows. The opposite case is strong and may prove right: large banks have more specialists, more pressure from supervisors and more volume over which to spread the cost.

What would undo the case

One development would remove most of the case: AI good enough at working out a bank's ontology from its existing systems that writing one by hand no longer pays. The limited evidence so far is against it. In one small 2025 study, of 20 questions on a single document, knowledge graphs built by automated methods answered 15 to 20 per cent correctly on their own. With the source text kept alongside them, they answered 90 per cent.⁶⁴ This is also the assumption most likely to change. The 2026 re-runs of the insurance test show part of that change: current models now answer most of its questions without written definitions. A bank committing to a five-year programme should test this assumption every year and build that review into the programme's governance.

Twice in recent years, Singapore's banks found they could not answer an important question on their own: whether a trade document was already being financed elsewhere, and whether a customer had raised concerns at another bank. Both times the answer was to agree what something is and write it down. Both arrangements took years to build, and neither needed AI. AI agents now make the same work necessary inside each bank, because they act on definitions at a speed and scale that leave little room for people to correct them. Banks that do this work will be able to use agents on terms their boards and supervisors can accept. Others are likely to find their agents confined to demonstrations.

⁶³ Global Legal Entity Identifier Foundation, 26 January 2026.

⁶⁴ da Cruz, Tavares and Belo (2025).

Appendix: The Terms in Plain Language

Agent, or AI agent Software that acts on a request instead of only answering it. It plans a short sequence of steps, uses other systems on the user's behalf and changes something at the end, such as releasing a payment, closing an alert or updating a record.

AI model, or large language model The technology behind current AI assistants. It predicts text, which makes it good at drafting and weak at questions that depend on following links between records. Not to be confused with the credit, market and capital models that a bank's model risk function oversees.

Benchmark A fixed set of questions or tasks with known correct answers, used to compare AI systems on equal terms.

Connection standard A common way for an AI agent to reach other systems, tools or agents, such as the Model Context Protocol or agent-to-agent protocols. It governs how an agent gets access, not what the data it reaches means.

Entity resolution Deciding whether two records refer to the same real-world party. It is usually done by scoring how alike the records look. In an ontology the decision becomes a recorded fact, with a date and an author, that can be reviewed and reversed.

FIBO The Financial Industry Business Ontology: a free, published description of financial concepts and their relationships, maintained by the EDM Council and updated quarterly. A starting vocabulary rather than a finished model.

Graph A network of connected records, such as clients linked to their parent companies and facilities linked to their collateral. Not a chart.

GraphRAG Retrieval-augmented generation in which the material given to the AI model comes from a knowledge graph rather than from loose passages of text, so the model receives the connections as well as the content.

Hallucination A fluent, confident and wrong answer from an AI model. In banking it is dangerous because nothing in the output shows that it is wrong.

Knowledge graph An ontology filled with an institution's own records, so that the links between a client, its parent, its facilities and its collateral are stored as facts rather than worked out again each time.

Ontology A written description of the things a business deals in, the relationships between them and the rules that follow, set out precisely enough for software to act on. It works like a dictionary that also states the rules.

Ontology graph An ontology together with the knowledge graph built on it.

Orchestration tool Software that manages the sequence of steps an AI agent takes and the hand-offs between agents, such as LangGraph, CrewAI or Microsoft Agent Framework. It coordinates planning, delegation and retries. It does not define what the bank's data means.

Query language The language in which a question is put to a database. SQL is the long-established one for tables; SPARQL and GQL are the equivalents for graphs. The AI model writes the query and the data supplies the answer.

Retrieval-augmented generation (RAG) The approach most banks tried first: split documents into passages, find the passages that most resemble the question, and give them to the AI model to answer from. It works for drafting and lookup, and fails when the passages that resemble a question are not the ones that answer it.

Rule, or axiom A statement in an ontology that produces further facts, such as: if A controls B and B controls C, then A controls C. Rules are what make an ontology more than a dictionary.

Schema The structure of a database: which tables exist and which columns they hold. It records how data is stored, not what it means to the business.

Semantic layer A defined set of business terms and calculations, such as exposure or active client, that sits between a database and the people or AI systems querying it, so that each term is calculated the same way every time. A narrower relative of an ontology.

Taxonomy A classification scheme: a hierarchy that says where something is filed. It supports counting and reporting, and says nothing about how the thing filed relates to anything else.

Vector, or embedding A list of numbers representing a passage of text, arranged so that passages on similar subjects sit near one another. It is what makes search by resemblance possible, and why that search is unreliable for questions of identity.

References

- Ai, Yi. 17 August 2026.** "Bounded Semantic Planning and Deterministic Compilation for Reliable Enterprise Text-to-SQL." arXiv:2608.16663. Preprint, not yet peer-reviewed; no link to data.world or dbt Labs is stated. Uses dbt Labs' adaptation of the ACME Insurance benchmark: 38 questions, three runs each. The governed approach answered 37 questions (97.4 per cent) correctly in all three runs, against 21 (55.3 per cent) for direct generation from the database schema. Across 114 runs the governed approach produced no wrong-but-executed result and one refusal; the baseline produced 29 wrong runs. <https://arxiv.org/abs/2608.16663>
- Allemang, Dean, and Juan F. Sequeda. 21 May 2024.** "Increasing the LLM Accuracy for Question Answering: Ontologies to the Rescue!" data.world, arXiv:2405.11706. Extends Sequeda, Allemang and Jacob (2023). Reports accuracy of 72.55 per cent once ontology-based checking and repair of each generated query are added; a further 8 per cent of questions returned "unknown", leaving an error rate of 19.44 per cent. <https://arxiv.org/abs/2405.11706>
- Artificial Analysis. Accessed 1 October 2026.** "τ³-Banking Benchmark Leaderboard." A 97-task version of the τ-Knowledge banking domain, scored by pass@1. On the date accessed the highest-scoring model completed 51.3 per cent of tasks and the next two about 50.5 to 51 per cent. <https://artificialanalysis.ai/evaluations/tau3-banking>
- AsiaOne. 4 March 2026.** "Hin Leong Founder O.K. Lim's Jail Term Cut to 13.5 Years after Appeal to High Court." Reports the High Court's decision of the same day reducing Lim Oon Kuin's sentence from seventeen and a half to thirteen and a half years; the convictions stood. <https://www.asiaone.com/singapore/hin-leong-founder-ok-lims-jail-term-cut-135-years-after-appeal-high-court>
- AsiaOne. July 2025.** "MAS Issues \$27.45m in Penalties to 9 Financial Institutions over \$3b Money Laundering Case." Reports that MAS identified the nine institutions through supervisory examinations conducted between early 2023 and early 2025, and that the breaches were connected with the 2023 money-laundering case. <https://www.asiaone.com/singapore/mas-issues-2745m-penalties-9-financial-institutions-over-3b-money-laundering-case>
- Association of Banks in Singapore. 23 June 2023.** "ABS Launches Trade Finance Registry to Counter Trade Finance Fraud." UOB was Chairman Bank, and UOB, Standard Chartered, DBS, Citibank N.A., OCBC and BNP Paribas were the project leads; some 30 other banks participated, including CIMB, Bank of China, Deutsche Bank, ING, MUFG, SMBC and State Bank of India. Banks register transactions of corporate customers, with the data hashed into an encrypted format, and duplicates trigger notifications in near real time. https://abs.org.sg/docs/library/abs-launches-trade-finance-registry-to-counter-trade-finance-fraud_v3.pdf
- Basel Committee on Banking Supervision. January 2013.** "Principles for Effective Risk Data Aggregation and Risk Reporting" (BCBS 239). Banks identified as global systemically important were to meet the Principles by January 2016. <https://www.bis.org/publ/bcbs239.htm>
- Basel Committee on Banking Supervision. 28 November 2023.** "Progress in Adopting the Principles for Effective Risk Data Aggregation and Risk Reporting." Of the 31 global systemically important banks assessed, only two were fully compliant with all the Principles. Cites fragmented IT landscapes, legacy systems and manual processes. Based on supervisory assessments up to 2022. <https://www.bis.org/bcbs/publ/d559.pdf>
- Bloomberg. 9 July 2020.** "After \$9 Billion Credit Hit, Banks Seek Trade Finance Revamp." Republished by the American Journal of Transportation. Reports that banks and other creditors faced "about \$9.26 billion in potential losses from these and other failures over the last six years." The named cases are Hin Leong Trading (about US\$3.5 billion), Agritrade International (US\$670 million), Hontop Energy (US\$473 million) and Zenrock Commodities Trading (US\$166 million), "all of them involving some form of alleged fraud, forgery, multiple cargo pledges or questionable accounting." The aggregate covers further failures beyond those four. <https://ajot.com/news/after-9-billion-credit-hit-banks-seek-trade-finance-revamp>
- Companies and Limited Liability Partnerships (Miscellaneous Amendments) Act 2024, Singapore. In force 16 June 2025.** Requires annual confirmation of registrable controllers' particulars and updates within seven days of confirmation; widens the definition of nominee shareholder; and requires nominee directors and nominee

shareholders to be filed centrally with ACRA, with existing entities to file by 31 December 2025. The stated aim is to reduce money-laundering risk by making nominee arrangements more transparent, in line with Financial Action Task Force standards. Only public agencies may access the full information maintained by ACRA. Summary by Rajah & Tann Asia: <https://www.rajahtannasia.com/viewpoints/amendments-to-enhance-transparency-of-beneficial-ownership-of-companies-and-llps-to-come-into-effect-on-16-june-2025/>

da Cruz, Tiago, Bernardo Tavares and Francisco Belo. 8 November 2025. “Ontology Learning and Knowledge Graph Construction: A Comparison of Approaches and Their Impact on RAG Performance.” arXiv:2511.05991. On 20 questions about a single document, graphs evaluated without their source text answered 15 to 20 per cent correctly, against 90 per cent for the same graphs with the text retained. <https://arxiv.org/html/2511.05991v1>

dbt Labs. 7 April 2026. “Semantic Layer vs. Text-to-SQL: 2026 Benchmark Update.” Re-runs the ACME Insurance benchmark created by Sequeda and colleagues at data.world: 11 questions, each run 20 times. On the original schema, text-to-SQL accuracy rose from 32.7 per cent in dbt’s 2023 run to 64.5 per cent with Claude Sonnet 4.6 and GPT-5.3 Codex, and the dbt Semantic Layer from 60.5 to 72.7 per cent. After three additional data models, text-to-SQL scored 90.0 per cent (Claude Sonnet 4.6) and 84.1 per cent (GPT-5.3 Codex), and the Semantic Layer 98.2 and 100 per cent. dbt reports that the Semantic Layer returns an error when it cannot answer, while text-to-SQL returns a plausible wrong number. dbt Labs sells the Semantic Layer. <https://docs.getdbt.com/blog/semantic-layer-vs-text-to-sql-2026>

Deprez, Bruno, Toon Vanderschueren, Bart Baesens, Tim Verdonck and Wouter Verbeke. 2024, revised 2025. “Network Analytics for Anti-Money Laundering: A Systematic Literature Review and Experimental Evaluation.” arXiv:2405.19383. Reviews 97 studies and finds that network analytics increases predictive power, while cautioning that graph neural networks need care under class imbalance and network topology, and that synthetic data can give overly optimistic results. <https://arxiv.org/abs/2405.19383>

EDM Council. “FIBO — Financial Industry Business Ontology.” Available in RDF/OWL, organised into modular domains, and reviewed and updated quarterly by a team of ontologists and financial-sector subject matter experts; published on GitHub under an MIT licence. Regulatory reporting consistency, for example for BCBS 239, is among its stated uses. <https://edmcouncil.org/frameworks/industry-models/fibo/> and <https://github.com/edmcouncil/fibo>

European Central Bank. 8 June 2026. “ECB Announces Main Milestones for Roll-out of Integrated Reporting Framework.” The framework will consolidate statistical reporting requirements for euro area banks into a single standardised system. A public consultation is planned for the second half of 2027, a one-year pilot from the second quarter of 2030, and official reporting from the second quarter of 2031. <https://www.ecb.europa.eu/press/pr/date/2026/html/ecb.pr260608~6766ec7154.en.html>

European Central Bank. “Banks’ Integrated Reporting Dictionary (BIRD).” Described as “a voluntary collaborative initiative bringing together central banks of the European System of Central Banks (ESCB) and commercial banks,” free of charge, in which “the concepts underlying reporting obligations are all identified and described once.” Not related to the BIRD text-to-SQL benchmark cited elsewhere. https://www.ecb.europa.eu/stats/ecb_statistics/reporting/bird/html/index.en.html

European Union. Regulation (EU) 2024/1689 (the Artificial Intelligence Act) and Regulation (EU) 2026/1744. Annex III classifies AI systems intended to evaluate the creditworthiness of natural persons or establish their credit score as high-risk, with obligations for data governance, technical documentation, record-keeping, transparency and human oversight. Regulation (EU) 2026/1744 (the Digital Omnibus on AI), published in the Official Journal on 24 July 2026, moved the application date for Annex III obligations from August 2026 to 2 December 2027.

Financial Conduct Authority and Bank of England. 7 January 2020. “Digital Regulatory Reporting: Pilot Phase 2 Viability Assessment.” Gathered cost data from eleven large firms (banks, specialist lenders and challenger banks) and 126 smaller firms across mortgage and derivatives reporting, and found that “the lack of standardisation of the definitions and descriptions of data by firms is a significant barrier.” Estimated an overall efficiency factor of 30 to 40 per cent. Modelled payback: a single reporting domain built on new central infrastructure does not break even within ten years; the same domain reusing existing industry standards pays back in three years; one domain aligned to planned regulatory change in four; three domains rolled out incrementally in seven. For large firms, mortgage

reporting alone carried operating costs of about £450,000 and change costs of about £700,000 a year, 5 to 10 per cent of total regulatory reporting. Concluded that standalone viability for one domain was likely to be low while the approach warranted continued investment, and recommended small incremental steps. The pilot was not put into production. <https://www.fca.org.uk/publication/discussion/digital-regulatory-reporting-pilot-phase-2-viability-assessment.pdf>

Gartner. 25 June 2025. “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027.” Press release, citing escalating cost, unclear business value and inadequate risk controls. A forecast rather than a measurement. <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>

Global Legal Entity Identifier Foundation. 26 January 2026. “The LEI in Numbers: Global Transparency and Digitalization Push Drives LEI Adoption in 2025.” More than 2.93 million active LEIs at the end of 2025, growth of 13.5 per cent during the year, and an overall renewal rate of 56.7 per cent in the fourth quarter (61.2 per cent in EU jurisdictions, 49.5 per cent outside them). <https://www.gleif.org/en/newsroom/blog/the-lei-in-numbers-global-transparency-and-digitalization-push-drives-lei-adoption-in-2025>

Global Trade Review (John Basquill). 19 August 2020. “Analysis: Hin Leong's ‘Vicious Cycle’ of Trade Finance Fraud.” Summarises the report of the interim judicial managers (PwC) filed with the Singapore courts in June 2020. In March 2020 a cargo of 780,000 barrels of ultra-low sulphur diesel aboard a single vessel supported three arrangements: a Crédit Agricole letter of credit for a sale to Unipeq; Société Générale financing, with letters of credit issued by Rabobank and DBS, for a sale-and-repurchase with Glencore; and Natixis and OCBC letters of credit for a Trafigura–Winson arrangement. The bills of lading bear “the same vessel name and aggregate cargo” and differ only in date. The report also found forged bank statements, bills of lading, sales contracts, invoices, deal settlement slips and other documents; inventory overstated “by at least US\$809mn” for the year ending 31 October 2019; and export contracts presented to 15 banks, in relation to 58 import letters of credit, “not supported by any underlying sale transaction.” <https://www.gtreview.com/news/asia/analysis-hin-leongs-vicious-cycle-of-trade-finance-fraud/>

Gruber, Thomas R. 1993. “A Translation Approach to Portable Ontology Specifications.” *Knowledge Acquisition* 5, no. 2: 199–220. The source of the standard definition of an ontology as “an explicit specification of a conceptualization.”

Lei, Fangyu, and colleagues. 2025. “Spider 2.0: Evaluating Language Models on Real-World Enterprise Text-to-SQL Workflows.” *ICLR 2025*, arXiv:2411.07763. 632 tasks on enterprise databases on BigQuery and Snowflake, many with more than 1,000 columns. An agent framework using o1-preview solved 21.3 per cent of tasks, against 91.2 per cent on Spider 1.0 and 73.0 per cent on BIRD. <https://arxiv.org/abs/2411.07763>

Li, Jinyang, and colleagues. 2023. “Can LLM Already Serve as a Database Interface? A Big Bench for Large-Scale Database Grounded Text-to-SQLs.” *NeurIPS 2023 Datasets and Benchmarks Track*, arXiv:2305.03111. BIRD contains 12,751 question–SQL pairs over 95 databases in 37 professional domains. On the development set, GPT-4's execution accuracy rose from 30.90 to 46.35 per cent when it was given annotated external knowledge (domain definitions, synonyms, numeric reasoning and value illustrations), and ChatGPT's from 24.05 to 37.22 per cent. BIRD's “financial” database matches the PKDD'99 Czech bank dataset in the CTU Relational Learning Repository. <https://arxiv.org/abs/2305.03111>

Malay Mail. 23 October 2025. “Ex-Citibank Manager from China Jailed in Singapore over Billion-Dollar Laundering Case.” A former Citibank relationship manager was sentenced to two years' imprisonment for holding S\$481,678 in laundered cash, forging bank documents and obstructing justice. <https://www.malaymail.com/news/singapore/2025/10/23/ex-citibank-manager-from-china-jailed-in-singapore-over-billion-dollar-laundering-case/195618>

MIT Project NANDA. July 2025. “The GenAI Divide: State of AI in Business 2025.” Based on a review of more than 300 public AI initiatives, interviews with 52 organisations and 153 survey responses. Reports that “95% of organizations are getting zero return” from generative AI and identifies the core barrier as learning: most systems do not retain feedback, adapt to context or improve over time. Widely reported from 18 August 2025; its headline figure has been disputed.

Monetary Authority of Singapore. 2018 and 2019. “Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector,” co-created with the industry in late 2018; and the Veritas Initiative, announced 13 November 2019, whose consortium phases have covered credit risk scoring, customer marketing, predictive underwriting and fraud detection (the last two in insurance), publishing assessment methods, case studies and open-source code.
<https://www.mas.gov.sg/schemes-and-initiatives/veritas>

Monetary Authority of Singapore. 9 May 2023. Speech by Minister of State Alvin Tan at the second reading of the Financial Services and Markets (Amendment) Bill 2023, published by MAS. States that “MAS will issue a directive to participant FIs detailing the threshold criteria” for each of the three risks, together with the associated list of red flags. <https://www.mas.gov.sg/news/speeches/2023/financial-services-and-markets-amendment-bill-2023>

Monetary Authority of Singapore. 1 April 2024. “MAS Launches COSMIC Platform to Strengthen the Financial System's Defence Against Money Laundering and Terrorism Financing.” Co-developed by MAS and six commercial banks (DBS, OCBC, UOB, Citibank, HSBC and Standard Chartered), covering three risks: misuse of legal persons, misuse of trade finance for illicit purposes, and proliferation financing. Information sharing is currently voluntary and takes place where a customer's profile or behaviour shows “objectively-defined indicators of suspicion, or ‘red flags’.” The Financial Services and Markets Act was amended in May 2023 to provide the legal basis.
<https://www.mas.gov.sg/news/media-releases/2024/mas-launches-cosmic-platform>

Monetary Authority of Singapore. 4 July 2025. “MAS Takes Regulatory Actions against 9 Financial Institutions for AML-Related Breaches.” Composition penalties totalling S\$27.45 million: Credit Suisse Singapore Branch S\$5.8 million; United Overseas Bank S\$5.6 million; UBS AG Singapore Branch S\$3.0 million; UOB Kay Hian S\$2.85 million; Citibank entities S\$2.6 million; Bank Julius Baer Singapore Branch S\$2.4 million; Blue Ocean Invest S\$2.4 million; Trident Trust Company (Singapore) S\$1.8 million; LGT Bank (Singapore) S\$1.0 million. All nine failed to corroborate significant aspects of customers' source of wealth in higher-risk situations; eight inadequately reviewed unusually large or inconsistent transactions; five had customer risk assessment methods insufficiently robust to classify clients correctly; two failed to take timely and sufficient risk mitigation after filing suspicious transaction reports. MAS observed that most of the institutions had established AML/CFT policies and controls, and that the breaches arose out of poor or inconsistent implementation of them.
<https://www.mas.gov.sg/regulation/enforcement/enforcement-actions/2025/mas-takes-regulatory-actions-against-9-financial-institutions-for-aml-related-breaches>

Monetary Authority of Singapore. 13 November 2025. “Consultation Paper on Proposed Guidelines on Artificial Intelligence Risk Management for Financial Institutions.” Consultation closed 31 January 2026. MAS stated that the Guidelines “will cover different AI applications and technologies, including Generative AI, as well as newer developments such as AI agents.” The proposals require accurate and up-to-date AI inventories, risk assessment across impact, complexity and reliance, and proportionate controls over data management, fairness, transparency and explainability, with the board and senior management accountable. A twelve-month transition is proposed from the date the Guidelines are issued; they had not been finalised by October 2026.
<https://www.mas.gov.sg/publications/consultations/2025/consultation-paper-on-guidelines-on-artificial-intelligence-risk-management>

Monetary Authority of Singapore. 17 March 2026. “MAS Issues Prohibition Orders against Former Relationship Managers, Mr Wang Qiming and Mr Liu Kai.” Prohibition orders of sixteen years (Wang, formerly of Citibank) and seven years (Liu, formerly of Julius Baer), on the ground that neither is a fit and proper person; reported by AsiaOne.
<https://www.mas.gov.sg/regulation/enforcement/enforcement-actions/2026/mas-issues-prohibition-orders-against-former-relationship-managers-wang-qiming-and-liu-kai>

Monetary Authority of Singapore. 4 May 2026. “MAS Collaborates with Banking Industry to Harness Artificial Intelligence in the Fight Against Financial Crime.” A proof-of-value on pre-emptive scam detection with five banks, the Government Technology Agency and the Singapore Police Force, building models on combined bank data. Account numbers are hashed so that only the contributing bank can identify them, and all data is deleted at the end of the exercise. The banks were not named. <https://www.mas.gov.sg/news/media-releases/2026/mas-collaborates-with-banking-industry-to-harness-ai-in-the-fight-against-financial-crime>

Monetary Authority of Singapore. 3 July 2026. “Safeguards for Agentic Finance at Runtime (SAFR).” White paper, version 1.0, developed under MAS's BuildFin.ai programme with Ant International, Circle, HSBC, J.P. Morgan Chase, Manulife, Mastercard, OCBC and Visa. Introduces “a governance checkpoint between every agent decision and its execution,” so that “no agentic action reaches execution without having been declared, authorised, and assessed.” States that it does not constitute regulatory guidance or supervisory expectations. <https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/ai-safr/safr.pdf>

Monetary Authority of Singapore. 5 August 2026. “Written Reply to Parliamentary Question on Agentic AI in Financial Services.” Reply by Deputy Prime Minister Gan Kim Yong, Chairman of MAS, to a question from Mariam Jaafar. As reported, it states that MAS's proposed Guidelines on AI Risk Management will apply to all AI uses by financial institutions, including agentic AI, and that MAS expects to finalise them soon. <https://www.mas.gov.sg/news/parliamentary-replies/2026/written-reply-to-parliamentary-question-on-agentic-ai-in-financial-services>

MotherDuck (Jacob Matson). 6 February 2026. “Your Data Model Is the Semantic Layer.” Tests Claude Opus 4.5, GPT-5.2 and Gemini 3 Flash on the 500-question BIRD Mini-Dev set across 11 databases averaging about seven tables. Adding descriptive comments raised accuracy by 1.1 percentage points on the training split and 0.2 on the held-out split, and by 8.7 points on the hardest database; it lowered accuracy by 2.9 points on the easiest. MotherDuck is a database company. <https://motherduck.com/blog/bird-bench-and-data-models/>

Neo4j. 2 January 2023. “How DBS Connect Knowledge with a Knowledge Quotient Framework.” Account of a fireside chat at GraphSummit Singapore 2022 with Anand Sundarraman, Senior Vice President, DBS Bank, describing a framework built to “represent some of the knowledge that we have in our bankers' heads.” It consolidated application, transaction, payment, credit bureau and ACRA data and applied pathfinding, clustering and community detection to product recommendation, supplier identification and concentration risk across directors, shareholders, industries and geography. This is a vendor's account of a conference talk, not published DBS architecture. <https://neo4j.com/blog/how-dbs-connect-knowledge-with-a-knowledge-quotient-framework/>

OCBC Bank. 7 November 2017. “OCBC Bank Is the First Singapore Bank to Tap Artificial Intelligence and Machine Learning to Combat Financial Crime.” In a proof of concept over one year of corporate banking transaction data, the system reduced alerts not requiring further review by 35 per cent and increased the accuracy of identifying suspicious transactions by more than four times. <https://www.ocbc.com/group/media/release/2017/ocbc-bank-ai-and-machine-learning-to-combat-financial-crime>

Payment Expert. 21 November 2025. “Swift's ISO 20022 Cutover: The End of MT and a 20-Year Promise.” Swift's coexistence period for cross-border payments ended on 22 November 2025, retiring legacy MT formats such as MT103 and MT202 in favour of ISO 20022, first published in 2004. The article describes richer, structured data as “exactly what regulators and scheme operators have been asking for in AML, sanctions and fraud controls.” <https://paymentexpert.com/2025/11/21/swifts-iso-20022-cutover-the-end-of-mt-and-a-20-year-promise/>

QA Financial (Michiel Willems). 22 September 2026. “Singapore's MAS Puts Cross-Bank AI Shield to the Test.” Identifies as the central practical challenge whether participating banks can “align different transaction formats, customer-risk indicators and fraud classifications sufficiently for models to make reliable comparisons,” and paraphrases MAS managing director Chia Der Jiun as saying that a model performing strongly at one bank may behave differently at another. MAS has not published how the alignment is done. <https://qa-financial.com/singapores-mas-puts-cross-bank-ai-shield-to-the-test/>

Reuters. 16 April 2020. “Singapore Oil Trader Hin Leong Owes \$3.85 Billion to Banks.” Reports, citing industry sources, liabilities of about US\$3.85 billion to 23 banks, with HSBC, DBS, OCBC, Bank of China, Société Générale, Standard Chartered and ABN Amro among the creditors. Later reporting put the exposure at about US\$3.5 billion. <https://uk.finance.yahoo.com/news/singapore-oil-trader-hin-leong-071311932.html>

Sequeda, Juan, Dean Allemang and Bryon Jacob. 2023. “A Benchmark to Understand the Role of Knowledge Graphs on Large Language Model's Accuracy for Question Answering on Enterprise SQL Databases.” arXiv:2311.07509; Proceedings of GRADES-NDA '24. One enterprise schema in the insurance domain, based on the OMG Property and Casualty data model, and one model, GPT-4 as available in 2023, used zero-shot. Writing SQL directly, the model

answered 16.7 per cent of questions correctly; writing SPARQL against a knowledge-graph representation of the same data, 54.2 per cent. The authors were at data.world, which sold a knowledge-graph-based data catalogue and was acquired by ServiceNow in 2025. The results have been superseded by later re-runs with newer models. Benchmark materials: <https://github.com/datadotworld/cwd-benchmark-data>. <https://arxiv.org/abs/2311.07509>

Shi, Quan, Alexandra Zytek, Pedram Razavi, Karthik Narasimhan and Victor Barres (Sierra; Princeton University). 4 March 2026. “ τ -Knowledge: Evaluating Conversational Agents over Unstructured Knowledge.” arXiv:2603.04370.

The banking domain has 97 tasks, roughly 700 interconnected knowledge documents and tools for account updates, and grades each task against the final state of the bank's database. The best configuration (GPT-5.2 with high reasoning and terminal-based search) achieved 25.52 per cent pass¹; no configuration exceeded 13.40 per cent pass⁴, the share of tasks completed correctly on all of four attempts. The authors attribute failures to retrieving the right documents from densely interlinked knowledge bases and reasoning over complex internal policies. Code and data are public in the τ^2 -bench repository under an MIT licence. <https://arxiv.org/abs/2603.04370>

Singapore Police Force. 18 November 2024. “Founder of Hin Leong Trading Sentenced and Jailed for Cheating and Instigating Forgery.” Lim Oon Kuin was convicted on 10 May 2024 on two charges of cheating and one of instigating forgery, in respect of US\$111,683,939 obtained from HSBC, having directed employees to forge documents presenting fictitious oil sales as genuine, and was sentenced to seventeen and a half years. https://www.police.gov.sg/media-hub/news/2024/20241118_founder_of_hin_leong_trading_sentenced_and_jailed_for_cheating_and_instigating_forgery

Singapore Police Force. 18 November 2024. “Surrender to the State of Assets Seized in Major Anti-Money Laundering Operation in August 2023.” The ten foreign nationals arrested in the operation were all convicted and sentenced to between 13 and 17 months. They surrendered about S\$944 million in assets; 15 of 17 other foreign nationals who had left Singapore surrendered a further S\$1.85 billion; and S\$144.9 million remained seized in respect of the other two. The total, close to S\$3 billion, covers 27 people, not only the ten charged. https://www.police.gov.sg/Media-Hub/News/2024/20241118_surrender_to_the_state_of_assets_seized_in_major_anti_money_laundering_operation

Singapore Police Force. 24 October 2025. “Former Bank Relationship Manager Linked to \$3 Billion Anti-Money Laundering Case Sentenced.” A former relationship manager at the Singapore branch of Bank Julius Baer & Co. Ltd. was sentenced to four months' imprisonment for using as genuine a forged document, having guided a client in forging a Chinese tax payment certificate and submitted it to the bank to open her account. https://www.police.gov.sg/Media-Hub/News/2025/10/20251024_former_bank_relationship_manager_linked_to_3_billion_anti_money_laundering_case_sentenced

The Star. 5 July 2025. “Singapore Regulator Fines Nine Institutions over RM9.95bil Money Laundering Case.” Reports MAS as saying that the penalties “took into account various factors, including the extent of the FI's exposure to the persons of interest.” <https://www.thestar.com.my/aseanplus/aseanplus-news/2025/07/05/singapore-regulator-fines-nine-institutions-over-rm995bil-money-laundering-case>

Trade Finance Global. 8 October 2020. “Singapore Launches the World's First Blockchain-Based Solution Aimed at Preventing Double Financing Fraud.” On the platform developed under the leadership of Standard Chartered and DBS, built on Hyperledger Fabric and hosted on Microsoft Azure, which lets a bank check whether another institution has already submitted a given title instrument for financing by “comparing encrypted data points across multiple blockchain nodes,” so that “data can be matched across each bank's ledger, without ever leaving their location.” <https://www.tradefinanceglobal.com/posts/singapore-launches-the-worlds-first-blockchain-based-solution-aimed-at-preventing-double-financing-fraud/>

Wretblad, Niklas, and colleagues. 2024. “Understanding the Effects of Noise in Text-to-SQL: An Examination of the BIRD-Bench Benchmark.” ACL 2024. Found noise in 52 of the 106 questions in BIRD's financial (bank) database, including incorrect reference answers in 20.7 per cent. Corrected data: <https://github.com/niklaswretblad/the-effects-of-noise-in-text-to-sql>

About

About the author

Dr. Joey Tan is an Associate Professor at the NUS Asian Institute of Digital Finance and a Visiting Associate Professor at the Luxembourg School of Business. He is a technology executive based in Singapore, with Managing Director-level experience across Asia-Pacific and EMEA, including country and regional leadership at Amazon Web Services and at Alibaba, where he worked in the Globalization Office on cross-border commerce and digital trade infrastructure, and earlier roles at Oracle, Microsoft and Borland. He teaches executive programmes on digital ecosystems, innovation management and generative AI business strategy, and was recognised as an IMDA Singapore Digital Leader in 2026. He is co-author of *Driving a Sustainable Business Culture* (World Scientific, 2026).

About the institute

The Asian Institute of Digital Finance was established by the Monetary Authority of Singapore, the National Research Foundation and the National University of Singapore to develop deep capabilities supporting the needs of digital financial services in Asia, and to provide thought leadership across education, research and entrepreneurship in digital finance. Its research spans artificial intelligence and machine learning, digital assets and ledger technology, digital finance platforms, green finance technology and next-generation financial services, alongside work on cyber security, fraud and anti-money laundering.

Disclaimer

The views expressed in this paper are the author's own and do not represent the official position of the Asian Institute of Digital Finance, the National University of Singapore, or any organisation with which the author is associated. Every source relied on is publicly available and is cited in the references. Institutions are named on the basis of their own published statements, regulatory actions on the public record, and filings made by courts, administrators or judicial managers. Nothing in this paper is an assessment of any firm's current systems, controls or conduct, and no inference about present practice should be drawn from an account of past events. Figures, dates and regulatory timelines in a fast-moving area should be verified against the primary source before being relied on.