

R3 Reports

Bridging the Gap Between Investment Banking Architecture and Distributed Ledgers

Martin Walker

A hand is shown in the foreground, reaching out towards a network of nodes and lines. The background features a wireframe map of the world and several circular icons containing currency symbols: Euro (€), Dollar (\$), Pound (£), Yen (¥), and Rupee (₹).

r3.

A hand reaching out to touch a glowing globe of financial symbols. The globe is composed of a wireframe mesh and is surrounded by various currency symbols (€, \$, ¥, £, ₩, ₪) and a network of white dots connected by thin lines, suggesting a global financial network or blockchain technology. The background is dark and blurred, showing what appears to be a cityscape at night.

Contents

R3 Research aims to deliver concise reports on DLT in business language for decision-makers and DLT hobbyists alike. The reports are written by experts in the space and are rooted in practical experience with the technology.

1. Introduction **2**
2. What do we mean by Trade Processing **3**
3. The New Approach **9**
4. Conclusion **15**

Disclaimer: These white papers are for general information and discussion only and shall not be copied or redistributed outside R3 membership. They are not a full analysis of the matters presented, are meant solely to provide general guidance and may not be relied upon as professional advice, and do not purport to represent the views of R3 Holdco LLC, its affiliates or any of the institutions that contributed to these white papers. The information in these white papers was posted with reasonable care and attention. However, it is possible that some information in these white papers is incomplete, incorrect, or inapplicable to particular circumstances or conditions. The contributors do not accept liability for direct or indirect losses resulting from using, relying or acting upon information in these white papers. These views are those of R3 Research and associated authors and do not necessarily reflect the views of R3 or R3's consortium members.

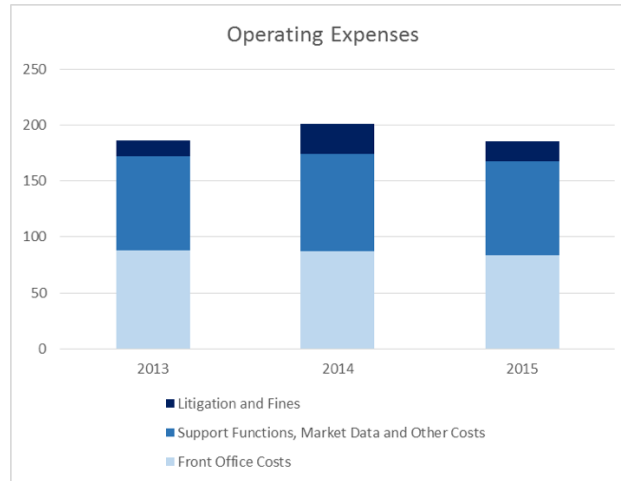


For more Research, please visit R3's Wiki [here](#).

1 Introduction

Specialist investment banks and the “markets” divisions of universal banks execute millions of trades each day, with total notional values running into trillions of dollars.¹ The banks’ trading partners are their customers (including fund managers, corporations and governments) and the major market counterparties (such as hedge funds and other banks).

This vast amount of trading activity is dependent on more than just traders. Inside each bank are complex IT infrastructures, and large numbers of support staff in critical functions such as Risk, Finance and Operations. Research by the Boston Consulting Group estimated the total cost of IT and other support functions at \$83.9 Billion in 2015. This number is both remarkably large and resistant to the banks’ cost reduction efforts, as shown in the chart below.²



Supporting all this trading activity (and generally reflected in the costs discussed) are many external organisations (and their infrastructures) including brokers, market data providers, central securities depositories (CSD), trading platforms, exchanges, matching platforms and clearing houses. The “front-to-back” processing of a trade typically will involve many systems (internal and external to the bank) and many parties.

Investment banking has always been a highly cyclical business, with profits dramatically rising and falling in line with market activity.³ Previously, one of the key shock absorbers of this volatile business was staffing costs. Good times meant large bonuses, bad times meant no bonus and staff cuts. However, since The Great Recession of 2007-2009, trading activity has generally been subdued, and normal operating costs have been hard to cut (despite significant reductions in front office size). Still, control and compliance costs have increased enormously, along with ongoing fines and litigation costs.⁴ Adding to these problems, regulations limiting bank trading activity, requiring additional capital, and setting targets for liquidity and leverage have severely reduced the opportunity to generate additional revenues.

It is unsurprising in such an environment that investment banks are more focused than ever on attempting to increase profitability by cutting operating costs. However, cost inflexibility means they have had to consider a range of new (and not so new ideas). These ideas include increased off-shoring, outsourcing, mutualisation of business functions, digitalisation and the application of distributed ledger technology (DLT).

Given the amount of interest in DLT over the last two years, it is worthwhile to pause and reflect. While many banks, consortia and fintech companies have done a great job in testing the technology, how will DLT help reduce those \$80-plus billion in operating costs? To encourage adoption of DLT in global markets, it is probably worth reflecting on some basic questions. What are fundamental needs for processing markets trades in a bank? Why have the associated costs

¹<http://www.bis.org/publ/rpfx16.htm>

²Source data available at <https://www.bcgperspectives.com/content/articles/financial-institutions-technology-digital-value-migration-global-capital-markets-2016/?chapter=3>

³Investment banking includes arms of the bank indirectly related to trading, such as equity issuance, debt issuance and M&A (Mergers and Acquisitions).

⁴Fines and litigation costs have been so large in both absolute and relative terms they have formed one of the main drags on profitability.

been so hard to cut? How far away are we from DLT nirvana, and what bumps are in the road (and which of those bumps are large enough to potentially force DLT off the road)? What could be done practically in the short-to medium-term to avoid some of those bumps and to assimilate key elements of DLT into the banks?

This paper argues that a pragmatic, hybrid approach is necessary to roll out DLT in capital markets, using elements of DLT and more established technologies (but applied in new ways). This assumes that even if DLT takes off on a large scale in investment banking it will be part of a mixed environment of centralised and distributed systems for a long period, and arguably forever. A pragmatic, hybrid approach would also make it easier for production-quality DLT solutions to meet the banking sector’s strict requirements regarding procurement, security and data privacy.

To create value for banks in the short-term DLT needs to be applied in a way that is not dependent on a “big bang” replacement of infrastructure. Even if significant DLT-based applications are rolled out, they will need to integrate with a great deal of existing infrastructure. Those who have worked in investment banking IT departments know that, despite the glamour of implementing new trading algorithms and low-latency trading infrastructure, the bulk of the work comes down to the unglamorous but critically important job of integrating systems – “the plumbing”.

This paper argues that a hybrid approach that includes elements of Big Data/Analytics with the “nudges” of behavioural economics can help banks tackle one of the major causes of their problems in implementing large-scale change: complexity. Complexity makes organisations (and systems) hard to measure, hard to understand and, consequently, hard to change. DLT, combined with Big Data, has the potential to replace the transparency and feedback loops that are missing in complex organisations. With those in place there is the potential to “nudge” organisations towards the standardisation and behavioural change, which would ultimately reduce costs and operational risk.

2 What do we mean by Trade Processing?

When looking at the complexity in an investment, it is very easy to lose sight of what all those systems and departments, spread across so many locations, are ultimately trying to achieve. It is important not to trivialise the amount that needs to be done, but it is possible to look at trade processing in way that clarifies things.

At its most basic level, trade processing within a bank consists of taking a small set of inputs from external systems or the physical world (i.e. traders) to generate a very large number of outputs. These are the outputs needed by a bank to answer fundamental questions such as, “How much money is being made?”, “What risks are being taken?”, “What trades could or should we do next?” There are also the outputs necessary to meet the requirements of regulators, accounting standards bodies and clients.

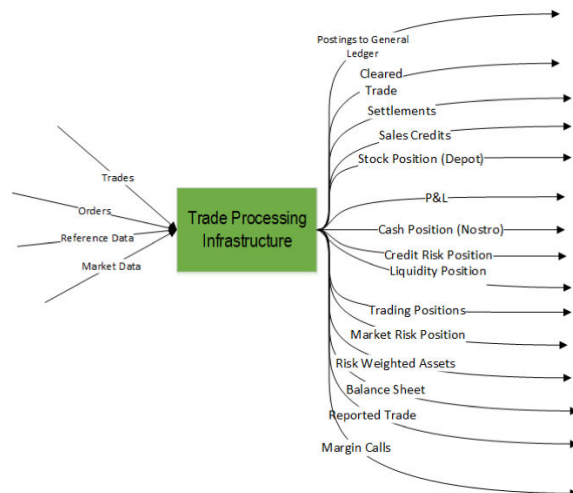


Figure 1: Inputs and outputs of trade processing

This does not, of course, mean that the optimal infrastructure would consist of simply creating

the green box as a single “smart contract”.⁵ There would be nothing particularly simple about turning all the required business logic into a DLT-style smart contract and it would be a very complex piece of code.

2.1 A model infrastructure

Another key point to understand about the trading infrastructure of most investment banks is that, broadly speaking, they work – trades get settled, credit limit utilisations get updated, postings get made to the general ledger. The extent to which they work in a cost-effective, flexible and controlled way, however, varies enormously between banks. Despite the problems, which to some extent are shared across all banks, it is possible to draw a picture of what “good” looks like, using existing technologies and techniques.

Looking across the trade processing stacks of most banks, for each main asset class, a pattern emerges. A low-latency, connectivity layer connects the trading businesses with exchanges, ECNs, SEFs and the firm’s own external-facing trading platforms. This needs to be both fast and highly resilient. Bursts of tens of thousands of trades interact with this infrastructure over the course of seconds. Those trades need to feed into the bank as smoothly as possible so they can be quickly credit-checked, executed, hedged and used to update positions along with bids and offers. Speed is essential because a slow-moving bank will find itself on the wrong end of trades with faster-moving rivals, or even faster algorithmic hedge funds. Resilience and lack of friction are just as important as speed. There is no point in being able to execute large numbers of trades and orders if they get “stuck in pipes” due to connectivity or static data issues.

While some elements of this low-latency infrastructure can be bought off-the-shelf, some banks have had very clever technologists building very clever infrastructure in this area. And don’t forget the very highly paid IT professionals who build this type of infrastructure in the hedge funds, which can often be far more sophisticated than the infrastructure of the banks they trade with.

Trades and orders are captured in this layer, and are typically fed to a trading system.⁶ The trading system was once the place where trades were booked, positions managed, risk and P&L calculated and viewed, i.e., where traders did their trading. Today, “trading system” is turning into something of a misnomer, with ever greater proportions of trades being executed on external platforms, more trading decisions being made by automated processes, and risk and P&L being viewed on a cross-asset class basis.

However, there is still the need for a central view of risk and P&L that human beings can view. There are also the more complex, the less liquid, and the “voice” trades, all of which need people and a trading system. Good trade infrastructure below the low-latency layer may not need to work at quite such a frantic pace, but it also requires being well-integrated with other infrastructure to have good quality reference data and the right business logic. Both internal and vendor systems in this space are often very mature (in a good way) and sophisticated. A bank can mess up the integration and configuration of even the best systems, but good systems combined with good integration can create very smoothly running infrastructure in this area.

It would be easy to state that trades then are fed from the trading system to the settlement system, where trades are settled. However, a better way of describing a settlement system (at least in a generic, cross-product way) is a place where trades are made ready for settlement. Cash flows and stock movements may be generated from the trades and trade events received from trading systems. Records of stock at depots and cash in nostro accounts may be updated and various exceptions resolved by Operations staff. When everything is ready, the instructions to move cash or securities are communicated to the outside world (or in some cases other parts of the bank).

The settlement infrastructure (and often closely related infrastructure for post-trade confirmation and matching) is typically the place where the noise of bad reference data and incorrectly booked trades become apparent. However, if reference data feeds are good, everything upstream is well-integrated, and the traders show a disciplined approach to booking their trades, the operational systems can work relatively smoothly.

Counterparties can still inflict pain on the best Operations department. They can mishandle their side of the trade, they can demand strange quirks in the post-trade processing, and, sim-

⁵A reasonably commonly accepted definition of a smart contract is code that can run autonomously to enforce and execute the terms of a contract. However, common usage of the term may vary between any programs running in a distributed ledger to genuinely independent objects that theoretically could have their own legal personality.

⁶Trades, quotes, orders and availability also generally are generated in this layer to feed to external platforms.

ply, they can be unresponsive when errors are found.⁷ However, a well-implemented post-trade infrastructure should be able to tell you the true cost of dealing with troublesome counterparties and provide the data to encourage better behaviour. Settlement systems share many of the same themes as trading systems. Many mature, high quality systems are available from vendors, but even the best of these can be disrupted by poor configuration and integration. Some banks still maintain their own settlement infrastructure. This may reflect the strength of having a mature high-volume system, though, in some cases, the persistence of in-house systems may simply reflect the difficulty of replacing them.

At the bottom of the trade processing layer are the systems that connect the bank to the places that will move the cash and/or securities, the custodians, CSDs, nostro banks, etc.

Interacting with the core trade processing systems are the systems owned by the major support functions, notably those associated with risk and finance. While the core trade processing systems are frequently specific to an asset class, those supporting other functions are generally cross-asset class. P&L needs to be generated for the bank and at the level of a trading desk or a book. Risk related to a counterparty needs to be viewed across all business lines. Though we have discussed at some length the trading and settlement systems, the costs of these systems can be much higher than the core trade processing systems, and they may contain a great deal of business logic.

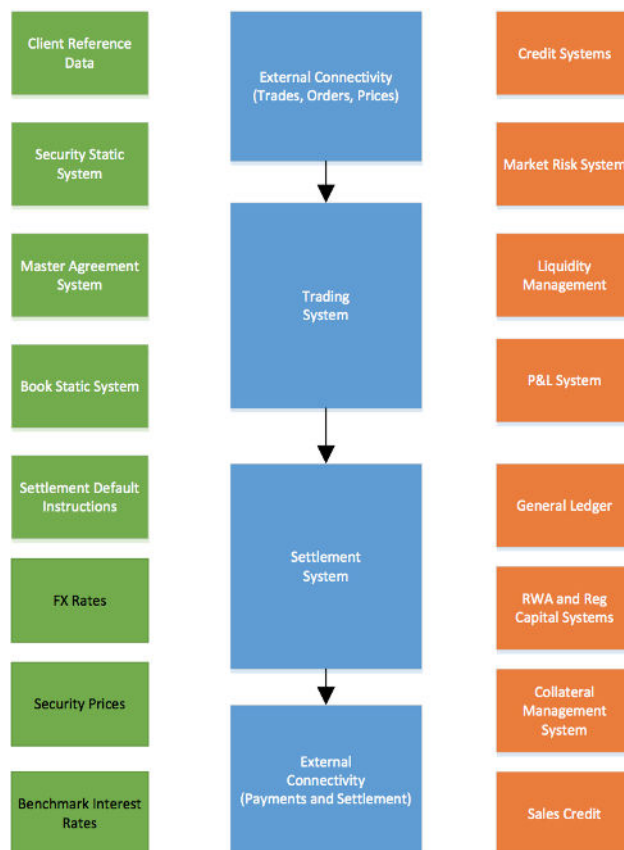


Figure 2: Core trade flow and related systems

Risk, finance and related systems have the same dependencies on good integration and good reference data. They work extremely well in some banks, but not so well in others, requiring significant manual intervention in core business processes. Perhaps the alarming consequence of poorly plumbed-in risk and finance infrastructure is that problems can remain obscured for much longer, as many banks discovered during the Great Recession.

Linked to the main trade processing systems (and, in many cases, the systems of support functions) are (or should be) centralized systems providing the reference and market data needed for calculations and trade enrichment.

⁷Examples include ad hoc changes to netted settlements, requests for non-standard information to be added to confirmation, and third party payments of the back of an FX transaction.

Acting as the backbone for most of this infrastructure is a messaging layer that allows communication between all these systems in near real-time and with guaranteed delivery.⁸ However, a significant degree of data still is exchanged as part of end-of-day batch processes, using a wide variety of methods. There may be SFTP, direct database connections, or even the use of the messaging layer. Whatever the methods used for communication, good infrastructure ensures the data reaches its destination in a timely manner without being mutated.

Other factors that should help the creation of smoothly running infrastructure are standardized messaging formats, such as FpML (for trades and the trade events), and a huge range of off-the-shelf software packages. These may be product-specific, function-specific, or a combination. An efficient and frictionless infrastructure can be created by adding all these elements together, using consistent market and static data, and making sure all traded products are fully supported by all the relevant systems, with messages that mean what they say.

2.2 The problem of trade processing

Investment banks today can process a vastly higher number of trades more quickly and at lower cost than in previous decades. Products that in previous decades could not even be imagined are now traded and processed daily. Banks particularly would like to reduce their technology costs, since they are high both in absolute terms and as a proportion of revenues.⁹ However, other sectors accept the need to spend money on technology. It may cost car manufacturers around a billion dollars to build a new plant, but they do not simply try to wish away the costs of doing business.

The difference in investment banking is the effort devoted frequently and fruitlessly to improve the infrastructure. Generally, the bigger and more ambitious the project, the greater the failure. There are front-to-back re-engineering of systems, “simplification”, horizontally-organised functional systems, clever (but expensive) middleware programmes. All these approaches have been tried multiple times across the banks with very variable rates of success. There have been many complete failures.

The other problem is the “noise” that builds due to progressive deviations from the model infrastructure described. Sources of this noise include:

- Problems in the quality and completeness of messaging between systems. Some front office IT departments use the terrible phrase “fire and forget.” Send the trade downstream, but do not worry whether anyone can make sense of it.
- Lack of reference data systems or failure to connect all relevant systems to those reference data systems.
- Poor trade booking front office staff, combined with a failure to encourage better practices.
- Trades and structures that are booked in one system, but not understood by the systems it feeds.
- Bespoke processes for clients.
- Significant volumes of manually booked “voice” trades that are inherently prone to error.
- Manual and/or paper-based confirmation and matching processes.
- Mutating standard messaging formats to make up for problems in other parts of the system infrastructure.

There are large variations in the degree of noise between banks, or within banks between different businesses or regions, but it is a theme almost every investment banker would recognise.

To put it more simply, superficially similar system infrastructures can have widely varying costs and levels of operational risk because some are not wired together properly, lack support for the products traded, or lack the relevant reference and trade data sourced from golden sources.

⁸The *messaging* system uses a built-in data store to persist messages It does not guarantee the recipient successfully can process the message

⁹Recent estimates by Accenture/McLagan for just IT costs (\$30 billion) were even higher than BCG estimates. https://www.accenture.com/us-en/_acnmedia/Accenture/Conversion-Assets/DotCom/Documents/Global/PDF/Consulting/Accenture-Banking-on-Blockchain.pdf

“Noise” or “friction” has other indirect costs in addition to the labour costs of paying people to fix problems. It needs more layers of control. In a fragmented, “noisy” infrastructure, those layers of bolted-on (rather than integrated) controls can become another source of noise and error. All the resulting complexity becomes progressively more expensive to run (and to change) because the quality of data required to make the right decisions deteriorates as it becomes more dependent on people and interpretation.

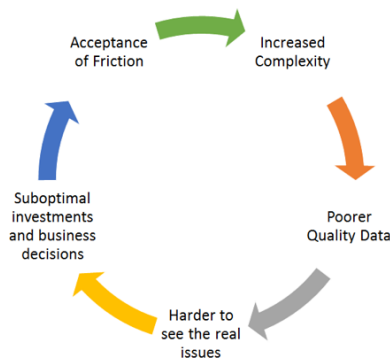


Figure 3: The complexity cycle

Regulators, the technical press and even senior bankers have grown highly critical of the state of investment banking infrastructure. Criticism extends beyond cost to operational risk, quality of data produced, and flexibility to deal with changing markets and regulations. One of the most recent tests of the banks’ flexibility to change was the trade reporting requirements under the Dodd-Frank Act (DFA) and the European Market Infrastructure Regulation (EMIR). These requirements were superficially straightforward: to report both the trade economics of OTC derivatives in near real-time and post trade data by the end of the day. The exercise proved problematic, if not traumatic, for most banks, with costs ranging from the tens of millions to hundreds of millions of dollars.¹⁰

All this complexity is not necessarily any specific person’s fault. Much of it has been the result of:

- Mergers and acquisitions by the banks, and failure to completely integrate infrastructure and businesses.
- Two decades of breakneck financial innovation, including the creation of hybrid products and structures.
- Decades of largely autonomous business units and legal entities making decisions that were right at the time for their entity or business, but ultimately not right for the organisation.
- The development of capital markets into a collection of genuinely global businesses.
- The errors that can arise every time a trade is translated from one system’s data structure into a messaging format, then translated again into the receiving system’s data structure.

2.3 The limits of centralisation and opportunity for DLT

One of the most positive areas in trade processing in recent decades has been the efforts to increase standardisation and centralisation¹¹ of processing. The huge increases in volumes and variety of products would not have been possible without:

- **Centralised** services, such as trading ECNs, the SWIFT network, CSDs, Euroclear and CLS.
- **Standardisation** of product definitions, legal documentation and message types through the work of organisations such as ISDA, SWIFT and ISLA.

¹⁰<http://www.bis.org/publ/othp20.pdf#page=41>

¹¹Not something you may have anticipated reading in a paper about distributed ledgers.

There has been an increased desire recently to extend this approach to the cost challenge. Banks have been more open than ever to the idea of mutualisation of their system functionality in areas that they do not consider to be differentiating. However, creating mutualised utilities has not proven easy. Infrastructure in the middle areas of trade lifecycle has proven particularly hard to mutualise because the systems in those areas must deal with the greatest amount of friction-driven complexity.

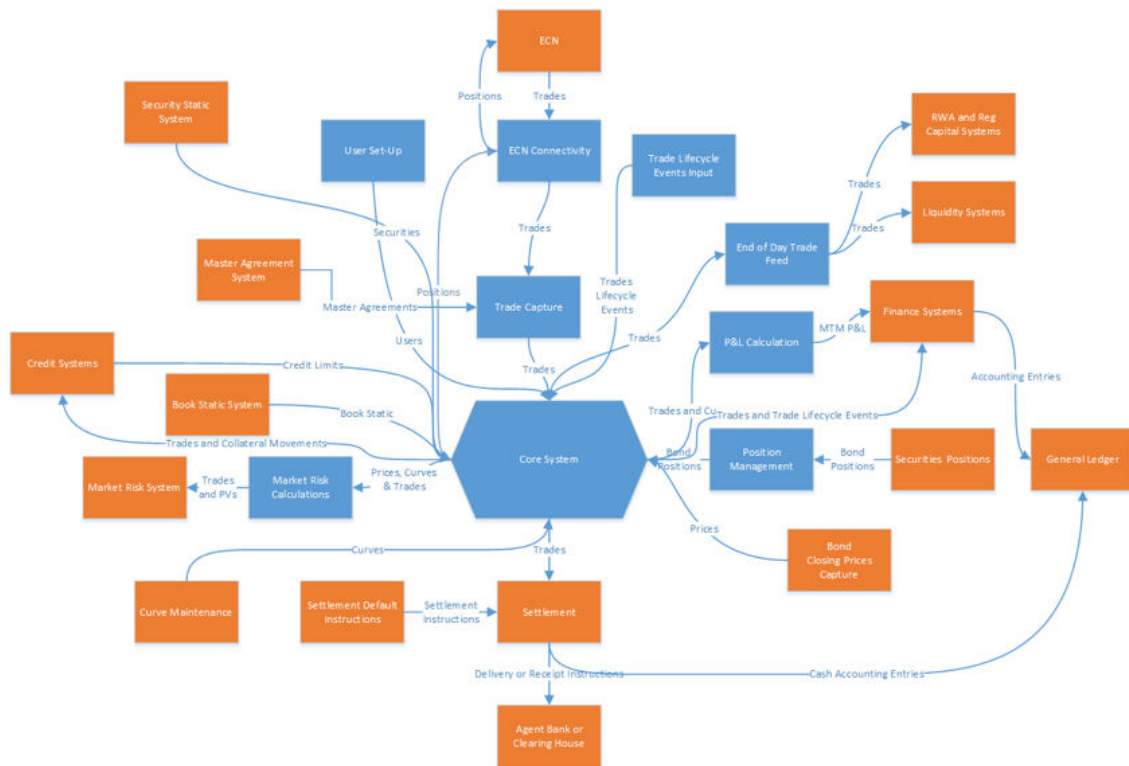


Figure 4: Repo trade processing (Simplified generic architecture)

Even for better built infrastructure, basic complexity theory kicks in because of the greater number of connections between systems in the middle of the trade processing lifecycle. To use technical terms, “complexity” is a function of the number of “edges” and “nodes,” as shown in the diagram above.

Ideally, a technology would be available that:

- has the same magic ingredient of the best centralised systems, i.e., standardised data models and business logic.
- deals with the basic drivers of complexity, i.e., reduces the number of edges.
- supports the mutualisation of non-differentiating processing between banks.
- can be combined with analytics software to make it easier to measure both problems and the impact of changes, providing a tool for better management decision making.
- potentially simplify the settlement process, and reduce the time needed for settlement cycles.

Many people would argue that DLT is that technology. However, how much of that is plausible in the short term?

2.4 The barriers to Ledger Nirvana in trade processing

In “Ledger Nirvana,” market counterparties use consistent sets of trade data with smart contracts (or CordApps in the Corda model) that apply consistent business logic to produce the various output required to operate a trading business.¹² However, despite many proofs of concept

¹²See Figure 1: Inputs and Outputs of Trade Processing.

in various aspects of capital markets, and the major investments made by some DLT-related companies, several obstacles must be worked through to enable more widespread of adoption of DLT in markets. *None are insurmountable, but all could take considerable time and effort.*

A ledger-based solution to trade settlement needs to be tangibly better in terms of cost, control, security, and resilience than existing financial market infrastructure, which generally works remarkably well. Examples include the major Central Securities Depositories and CLS Bank in the FX Market.

The Ledger Nirvana settlement infrastructure, is typically based on the assumption of cash and securities “on ledger”. In other words, cash or securities legally exist on the ledger, or have a “tokenised” representation of assets that enjoys the same degree of legal certainty and settlement finality as the primary record of the assets. Some progress has been made in this area with Overstock’s issuance of shares on its T0 platform, which includes the use of a private ledger doing the core processing, while recording the transactions on the Bitcoin Blockchain. This is a step forward legally and technically, but is still a long way from making a significant impact.

Genuine “cash on ledger” is even more problematic. Fundamentally, real-world fiat currency needs to be a ledger version of central bank reserves (which requires the co-operation of central banks and legal changes, depending on the jurisdiction) or commercial bank-issued money. Commercial bank-issued money would be the direct economic and legal equivalent of the money that is represented today as positive balances in customers’ bank accounts. Direct equivalence, however, brings the challenges faced in the existing world. Commercial bank money carries credit risk against the issuing bank and needs a mechanism (equivalent to the central bank clearing systems) to control the credit risk that builds up between clearing banks as funds are transferred.

For securities such as government bonds and more liquid equities, mechanisms would need to be built to allow them to not just be bought and sold but to be provided as collateral between banks, to CCPs, and to central banks through their repo process.

A smart contract that carries out the mechanics of a financial product, such as the work done by Axoni on Equity Swaps or Barclays and R3 on Interest Rate Swaps, represents a step forward in supporting the trade lifecycle of derivatives trades using DLT. However, there are many systems within the trading infrastructure of a commercial bank that execute, enrich, process and aggregate trades and trade events. A smart contract that performs the basic mechanics still would need to interact with credit risk, market risk, liquidity management, position viewing, P&L calculation/aggregation, regulatory reporting, derivative clearing, sales credits and many other systems. What is frequently forgotten is that simply having a ledger of trades does remove the need for what is frequently the most complicated and expensive system in a bank, the general ledger. A typical general ledger system is not just a list of transactions. It also is a list of accounting rules and policies that are applied to the transactions, often requiring the support and judgement of a large finance department.

In many markets such as spot FX, futures, equities and the more liquid bond issues, most trading (including much order processing) takes place at very high speed using very expensive and sophisticated infrastructure. It can be argued that this speed does not add significant value to society or the economy, but that is the way many markets operate. There would be great resistance to any attempt to slow down trading to allow DLT, which is inherently slower, to replace the current pre-trade infrastructure. A further obstacle is that the post-trading processing costs of electronically executed trades are considerably lower than for the more traditional (and error-prone) voice trading. This means DLT solutions need to be a significantly better post-trade solution for the overall volumes of trades in an asset class.

Ledger Nirvana also would position distributed ledgers in markets as Financial Market Infrastructure within the scope of the Bank for International Settlements “principles for financial market infrastructures.”¹³<http://www.bis.org/cpmi/publ/d101a.pdf> These principles are incorporated in law in most markets, and are very demanding. They represent a high, but necessary, hurdle for DLT to clear.

Overlapping with the BIS principles are the banks’ own requirements for high-volume processing, resilience and security. Just as the existence of a distributed ledger does not automatically remove the need for a general ledger, the use of cryptographic techniques does not make a system secure from a bank’s perspective.

In Ledger Nirvana, the trade is the settlement. A trade is booked and value is exchanged. However, this creates significant problems for today’s business models, business models that cannot

¹³<http://www.bis.org/cpmi/publ/d101a.pdf>

simply be wished away by the DLT enthusiast. Most of capital markets work implicitly on time delays. Huge daily volumes are traded and processed, but a market maker only needs to be flat (in most markets) by the end of the day. The settlements teams only need to transfer the net settlement amounts at the end of the settlement cycle. In the world of “trade equals settlement,” a market maker can only create liquidity for the market in one of two logical ways.

1. They can “warehouse,” i.e., stockpile what they are buying and selling. Under current regulations this incurs capital charges that would make market making completely uneconomical.
2. If they do not warehouse, selling by a market maker would require a mechanism for near instantaneous borrowing of securities and buying. Buying would require either a large credit facility or near instantaneous financing of the bought assets.

Overall, these barriers could delay “Ledger Nirvana” by years.

3 The new approach

Given the opportunities to improve trade processing, but noting the barriers to large-scale adoption of DLT, is there scope for an intermediate/hybrid approach that uses some of the elements of DLT to focus on the specific causes of problems identified above?

- Assumptions:
 - Within and between banks, data business logic is already highly distributed, but it is prone to the creation of friction in many parts of trade processing.
 - Complexity mutualization or centralisation of the existing distributed logic and data hard, expensive and risky.
 - Work towards DLT nirvana will be slow and will possibly never reach the destination enthusiasts hope for.
- The way forward requires:
 - A new model for sharing data data models and business logic.
 - Mechanisms are needed for driving changes in behavior i.e. making the sources of cost and operational risk more transparent, just as the mining concept in Bitcoin created behavioural incentives.
 - A recognition that outputs of existing infrastructure (including their idiosyncrasies and noise) need to be captured by a data layer for post-execution processing.
- Solution requires:
 - DLT capabilities for reaching consensus, immutability, distribution of data and sharing of business logic.
 - Analytics and other mechanisms for behavioural change. Relating costs and operational risk back to specific causal factors.
 - Application of best of breed systems for capturing data, transforming it, and dealing with exceptions.

A high-level model follows that uses Corda (the R3 Distributed Ledger), analytics techniques, and existing bank standard ETL tools to meet the solution requirements.

3.1 Overview

The principle elements of the hybrid solution are:

- DLT can support a mechanism that allows banks to agree on how different products will be processed. For each financial product, a *Product Definition Agreement* will list the agreed formats of data and the collection of services and systems that will perform the relevant parts of the trade lifecycle.

- Existing technologies would be used to load incoming messages, validate them, and either create a new object (in accordance with the Product Definition Agreement) or link to existing objects.
- It will attempt to “link” incoming notifications from the other parties to trades (or related objects) currently stored, based on trade economics. (It does not wait for all trade attributes to create a perfect match.)
- Both sides of a linked trade are stored in the same data object (the Golden Container) and any updates (except for private data) are distributed to all relevant parties.
- A *Service Notification* process determines whether an object has reached a sufficient degree of completeness or consensus between parties to hand it off to services (whether smart contracts, existing bank systems or market infrastructure) that perform parts of trade processing. Conversely, if the degree of completeness or consensus is broken, it also will inform the relevant services.
- The key data for analytics tools covering cost per trade, operational risk and client efficiency will be provided by recording the capture of the trades and events, together with any exceptions, and the length of time spent processing.

3.2 Product definition agreement

The Product Definition Agreement provides the basic set of rules parties need to agree upon about how to process their mutual trades on a collaborative, distributed basis.

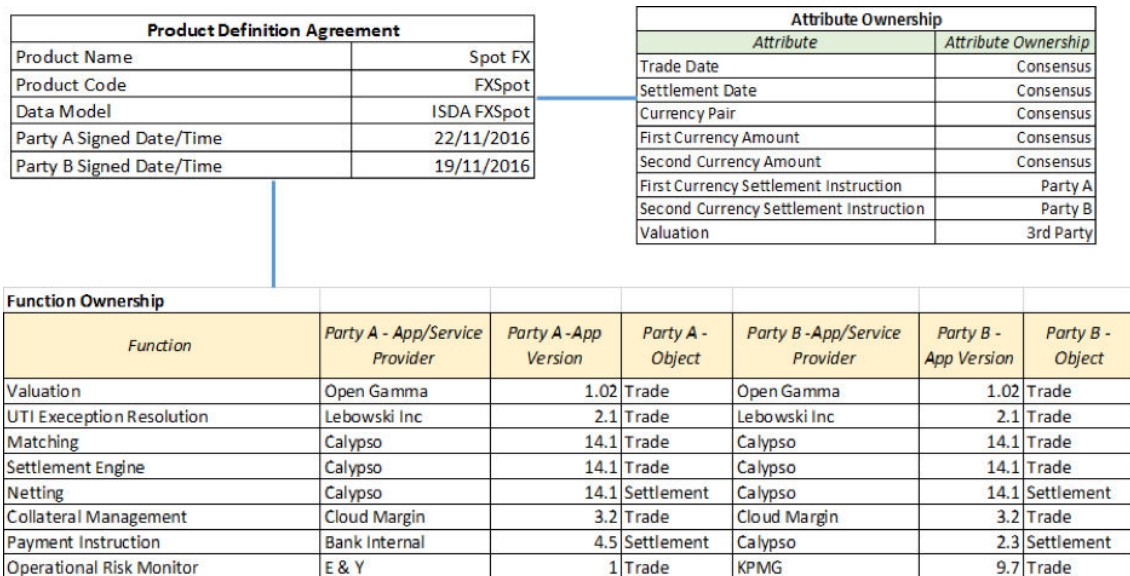


Figure 5: Sample produce definition agreement

The consensus mechanism in DLT, such as Corda, allows all relevant parties to “sign” that they are agreed to the processing rules and data structure.

The relevant parties can explicitly agree on:

- The structure of data for trades and events processed.
- The ownership of data attributes. For some data attributes the agreement may state that one party fully owns an attribute. One part may “own” the population of their own settlement instructions, which then are accepted by all relevant parties for use in the settlement process. For others, such as a valuation, a third party may own the right and obligation to populate the data. Finally, many attributes, such as trade economics, will be owned by both trading counterparties, and the attribute is only recognised as correct when both parties sign the agreement.

- **Functional Ownership** – it can be agreed that specific parts of the trade lifecycle can be carried out by a specific system or party. Both parties may agree to use a specific smart contract, a cloud version of an existing vendor system, a centralised service, or different systems (accepting the risk that this increases the chance of differences).

3.3 The Golden Container

Trade and event data is stored in what we describe as the “Golden Container”. There has been a seeming endless quest by many in banks for a “Golden Trade” record. The general idea is creating a single version of the truth for a trade (either at an organisational or a market level). Unfortunately, the idea tends to break down unless it involves a centralised system to not only store the golden record, but do all the processing. If multiple systems carry out different elements of trade processing, including lifecycle events, enrichment, or generation of derived data (settlements, aggregated positions etc.), there are many opportunities for the Golden Trade to become tarnished, even for the simplest cash products.

The concept of the “Golden Container” is very different from many DLT Markets proofs of concept that assume either consistency of the trade from outset or assume that the trade is either in an agreed or not agreed state. It still would use many of the strengths of DLT, but would drive towards a consistent view of the trades through the rapid identification of inconsistencies between both the parties involved and their systems.

The object model attempts to balance control and flexibility. Consistency starts by linking the two parties’ views of world, providing a path to consensus, early identification of differences, and views on the state of a trade (or related object) from multiple functional perspectives.

The most important step is creating a linked version of the trade that contains both parties’ views. Each time a relevant system updates a trade, it updates their view of the trade, and sends that update to their counterparty (or any other agreed relevant party).

The key benefit is that both parties can see in real time if they do something to make the trades diverge.

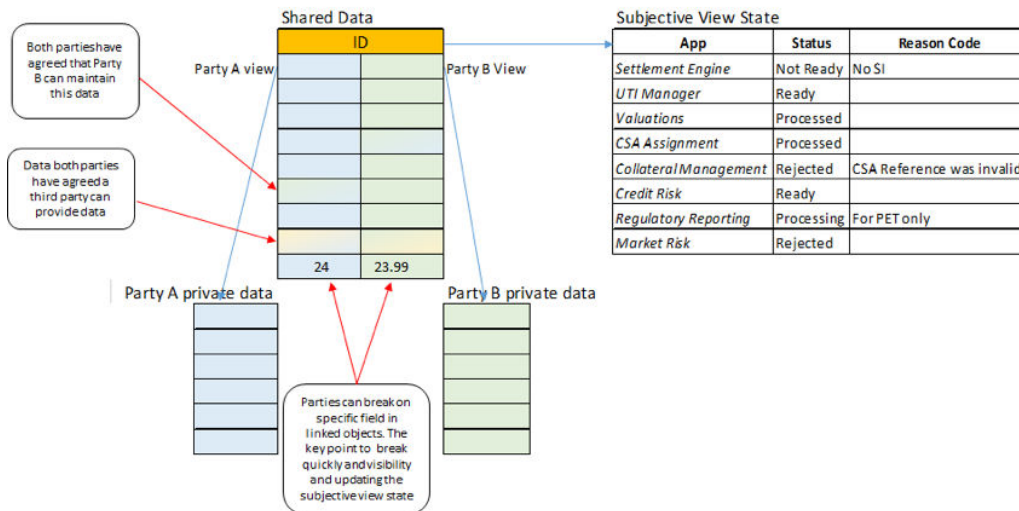


Figure 6: The "Golden Container"

The Golden Container contains two other key elements.

- **Private data** related to the trade will be stored only on the relevant party’s node. Private data includes data that genuinely needs to be kept private (P&L, trading book, etc.) and data that is only needed by one party, and should be kept segregated from the main trade object. This avoids mutation of the trade object through the addition of data superfluous to one party or the “over-loading” of fields.
- **Subjective View States** – the ability to validate and process objects from the perspective of different functions. For instance, the trade may be considered valid (and ready for further processing) by one application (e. g. Front Office, which requires only a calculated PV),

but may not be valid from another function’s perspective (e. g. collateral management may require that the relevant master agreement is assigned to the trade).

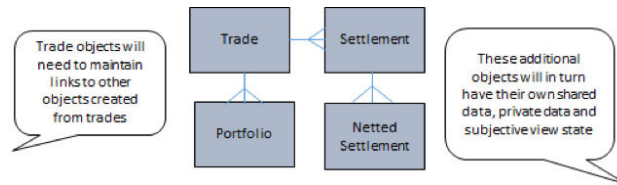


Figure 7: Relationship between key objects

3.4 Data capture

A key element of the model is data capture. Data needs to be captured within and outside the bank for any relevant trade or trade event.

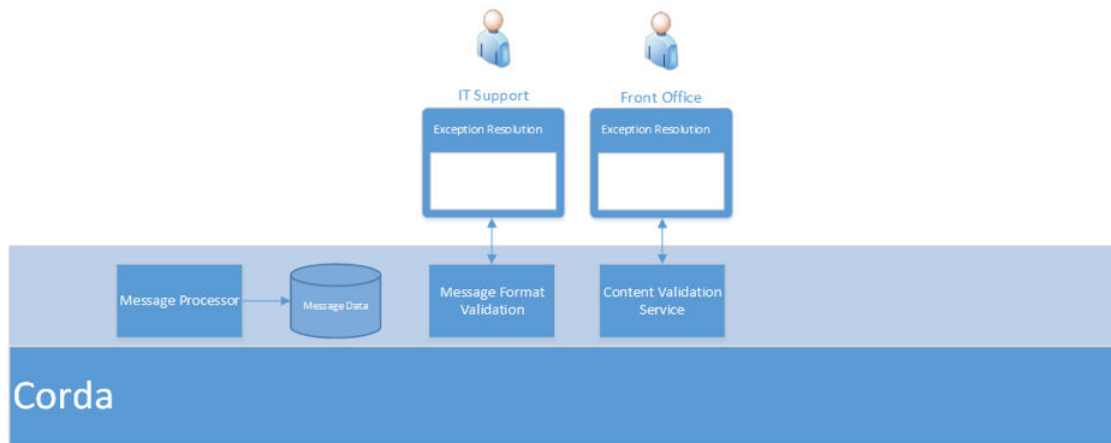


Figure 8: Data capture

The key first step to adopting a more distributed approach to trade processing is recognising the need to load trades (and other messages) into the ledger from existing sources. This could be a feed from a source where the parties already have agreed on the trade economics, such as an ECN, or from an existing trading system where trades are inputted manually and unilaterally by one side’s front office.

It must be able to process any popular message format used today. To maximise flexibility, trading counterparties should be able to agree to use a third-party message validating/translation service (see Product Definition Agreement). Several proven products on the market allow the data capture and data processing tasks to be carried out, largely as a configuration task rather than a major programming exercise.

These tools would allow the processing of existing message formats used by banks, including standard formats (e.g. FIX, FpML), company-specific formats, and bank-specific modifications of standard formats.

3.5 Distribution

Corda Nodes within the firewall of each bank would control the flow of messaging between the banks’ internal systems and, where relevant, to the nodes of other parties relevant to each trade. The actual secure transport layer could even be provided by an existing supplier of secure messaging that is already integrated into each bank’s infrastructure.

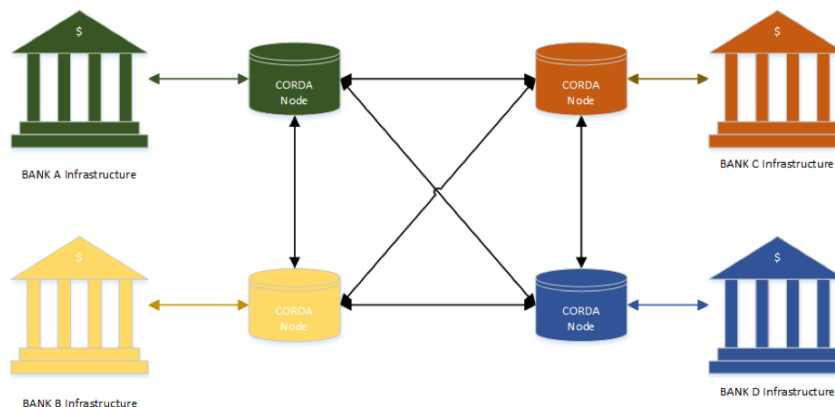


Figure 9: Data distribution

3.6 The status monitor

The structure of this model makes it straightforward to implement a status monitor that would allow a centralised support function or middle office staff to have a near real time view of the status of a trade (and related objects) from all perspectives.

Some banks have attempted to create something similar because of the potential benefits from a control and operational efficiency perspective, but they generally are held back by the fragmented state of their architecture.

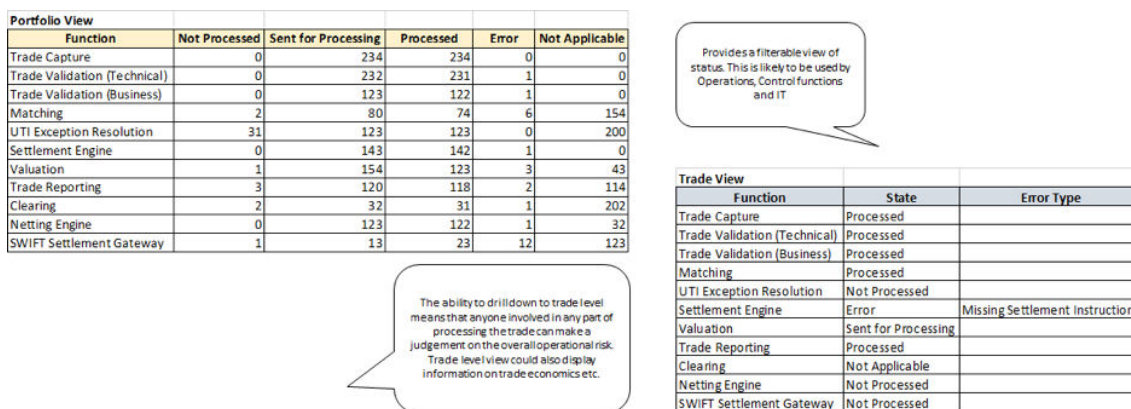


Figure 10: Example views of status monitor

A generic trade status monitor could revolutionise the management of operational risk in the trade lifecycle and potentially allow a greater deal of standardisation of process and error resolution in back and middle office teams.

3.7 Functional model

Overall, this design depends largely on the core capabilities of Corda and Extract Transform & Load (ETL) tools already in use in many banks. The new functionality consists of sets of small and relatively simple components.

The *Object Constructor/Updater* determines whether a new object, e.g. a trade or a settlement, should be created, or whether an existing object should be updated.

The *Linking/Matching Service* determines whether the new object can be linked to an object created based on data from a counterparty (or other third party). If it finds a sufficient degree of consensus, it will merge the objects into a single *Golden Container* (see below) that has both parties' versions of the object in a single data structure. There also will be the option for both parties to "force match" two objects into a single Golden Container, if both parties agree the trades are the same trade, even if they potentially disagree at that point about some of the details.

A *Service Notifier* will inform all relevant connected apps or services if the trade (or other object) has reached the state of consensus or completeness that allows another part of the trade lifecycle to take place, e.g., settlement or trade reporting. The Service “Notifier” also will inform the relevant services if an update’s object has changed sufficiently for a cancellation or amendment process to take place.

Product Definition Creator takes new product definitions (or amendments) and manages the process of getting the agreement signed by both parties and translated into the rules that the Service Notifier uses to communicate with services/apps.

Putting these relatively straightforward components (largely based on existing technology) together creates a platform that can be introduced within existing infrastructure.

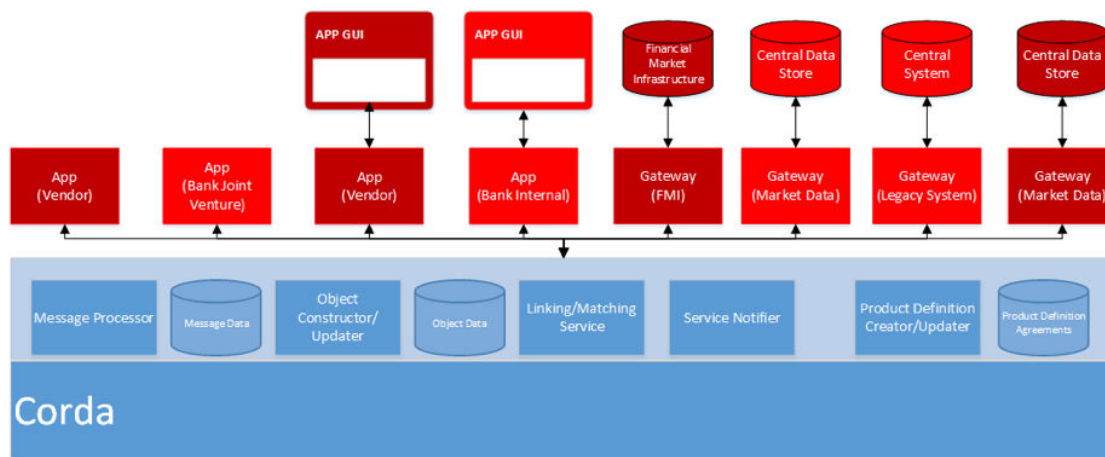


Figure 11: The overall ecosystem

3.8 The feedback loops - cost and control

Much of the money invested in IT still goes into making things faster rather than more transparent. The Economist, 2009

Last, but not least, are the “feedback loops.” These will make problems more transparent, potentially cut the cost of control, and provide the basis for better management decision-making by allowing the objective measurement of system or process improvements.

Generic tools can be created using the basic data contained within the rules engine (derived from the Product Definition Agreement) about the “meta-workflow” of trades, e.g., when it was matched, when it was complete enough for settlement, and the time delays in processing.

These tools (that could be produced in partnerships with organisations specialising in control and/or analytics, such as Audit firms) include:

Population-Based Audit Testing

Currently, the external audit process is based on samples, and implicitly does not trust the bank’s internal records, an immutable source of trade data that can be analysed in near real time or on large data sets.

Post-Trade Cost Analytics

The main drivers of the human costs in trade processing are exceptions to straight-through processing, and delays at any point in trade processing. The platform will collect the exceptions related to trades from function-specific processing apps, and allow more accurate cost estimate down to the level of the individual trade.

Quantitative Operational Risk

Measurement of operational risk is currently a highly manual, largely qualitative process, but one that has a major impact on calculated capital charges. The platform allows a quantitative measurement based on exception rates and time delays in processing.

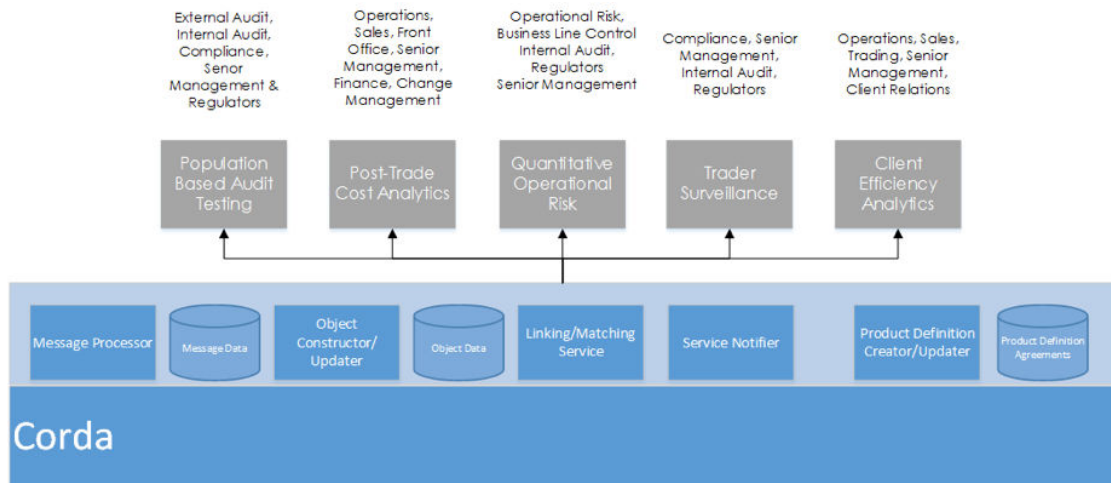


Figure 12: Integrating analytic components

Trader Surveillance

Current systems to detect rogue traders typically try to look for patterns in the overall state of a trade across all the relevant trade processing functions. An approach based on platform status monitor allows consistent business logic to be used by banks to look for anomalous patterns.

Client Efficiency Analytics

Many banks currently carry out “Client Efficiency” analysis to determine the relative cost of doing business with different clients, using the data to change pricing or drive change. The platform provides the opportunity to collect all the relevant data from a single place and rollout out the same analytics tool to multiple banks.

4 Conclusion

This paper gives a clear explanation of the mechanisms causing problems in the trade lifecycle processing and the techniques that could be used to address them. Many problems clearly are not technical, but socio-economic or political problems within banks. Nobody sets out to design bad infrastructure, but a succession of decisions, which are optimal for one specific function or business line, progressively can create complexity and drive entropy within the infrastructure and the organisation. Complexity reduces understanding of the systems, which drives more sub-optimal decision-making, creating more complexity.

One of the key factors driving this model is a recognition that in markets, agreement about the “facts” of a transaction (whether between the parties involved, or between systems within the same bank) can be highly unstable and asymmetric. Trade processing, in many areas of markets, is not simply about agreeing on the details and settling the trade. The trade may undergo many lifecycle events, and there may be changes to many of the non-economic attributes during the life of the trade. The view of trades will be partially asymmetric between parties because different banks have different risk appetites, different accounting treatments, or simply want to conceal information about the trade that the other party “does not need to know,” such as their own P&L arising from the trade. The more stable and symmetrical trade types, i.e., the simpler products, such as spot FX or cash equities trades, are likely to converge to the “purest” DLT model more quickly, but even they are likely to benefit from the feedback loops outlined in this paper.

Some DLT enthusiasts may argue there is relatively little DLT in this hybrid model, but that is by design. DLT as a technology will rise or fall based on its effectiveness in solving problems. This model introduces key benefits of DLT into the heart of trade processing in a relatively undistruptive way. It also provides the feedback loop that can objectively measure the success of different approaches, whether they use DLT, current technologies, or simply involve organisational or process change.



r3 is an enterprise software firm using distributed ledger technology to build the next generation of financial services infrastructure.

R3's member base comprises over 80 global financial institutions and regulators on six continents. It is the largest collaborative consortium of its kind in financial markets.

Consortium members have access to insights from projects, research, regulatory outreach, and professional services.

Our team is made of financial industry veterans, technologists, and new tech entrepreneurs, bringing together expertise from electronic financial markets, cryptography and digital currencies.

c·rda is an open source, financial grade distributed ledger that records, manages and executes institutions' financial agreements in perfect synchrony with their peers.

Corda is the only distributed ledger platform designed from the ground up to address the specific needs of the financial services industry, and is the result of over a year of close collaboration between R3 and its consortium of over 80 of the world's leading banks and financial institutions.