AUTOMATION

Bridging the gap between
investment banking infrastructure
and distributed ledgers

MARTIN WALKER I ANTON SEMENOV

AUTOMATION

#46
11.2017

THE HENLEY MBA

Realise your potential with an MBA
from a triple-accredited business school.

hly.ac/the-henley-mba

Henley
Business School
UNIVERSITY OF READING

Where business comes to life

JOURNAL

CONTENTS

BUSINESS MODELS

INVESTMENTS

Bridging the gap between investment banking infrastructure and distributed ledgers

MARTIN WALKER | Banking & Finance Director, Center for Evidence-Based Management

ANTON SEMENOV | Principal Business Analyst, Commerzbank AG

ABSTRACT

Enthusiasm for blockchain, or the broader family of distributed ledger technologies (DLT), within capital markets is now into its third year. The enthusiasm has manifested itself in numerous pilots, proofs of concept, fintech startups and industry collaborations. Within investment banks, large broker-dealers, and many "buy-side" firms, the enthusiasm has in large part been driven by a combination of heavy demands on IT departments and considerable pressures to cut costs. This paper argues that there is a key factor that has prevented the delivery of any significant working systems within these enterprises; a general tendency to try to fit a solution (blockchain and its derivatives) to problems, rather than trying to understand problems and find the appropriate solutions.

Relative analysis of firm infrastructure suggests that the root causes of most problems are not technical but are human in nature; being related to incentives, culture and organizational structure. The analysis demonstrates that trade processing data and business logic are highly distributed but frequently highly inconsistent. A model is proposed (drawing on many of the concepts of DLT) to create both transparency of issues and mechanisms for the propagation of consistent business logic and consistent data models. This aims to use DLT based techniques to deal with fundamentally human problems by the introduction of the appropriate feedback loops for management decision making. Understanding the nature of problems and the effectiveness of changes would allow genuinely evidence-based management decision making. A technologically driven, human transformation that could act as a lever for unravelling organizational complexity.

1. INTRODUCTION

Specialist investment banks and the "markets" divisions of the universal banks execute millions of trades each day, with total notional values running into trillions of dollars. The banks' main trading partners include fund managers, pension funds, hedge funds, large non-financial corporations, and other banks.

This vast amount of trading activity is not just dependent on traders. There are complex IT infrastructures inside each bank and large numbers of support staff in critical functions, such as risk, finance, and operations. Research by the Boston Consulting Group estimated that the total cost of IT and other support functions within major financial institutions that supported trading in financial products was around U.S.$83.9 billion in 2015. This is a number that is both incredibly large and remarkably resistant to the cost reduction efforts of banks, as shown in Figure 1.

Figure 1: Investment banking operating expenses (U.S.$ bln)

Front office costs

Support functions, market data, and other costs

Litigation and fines
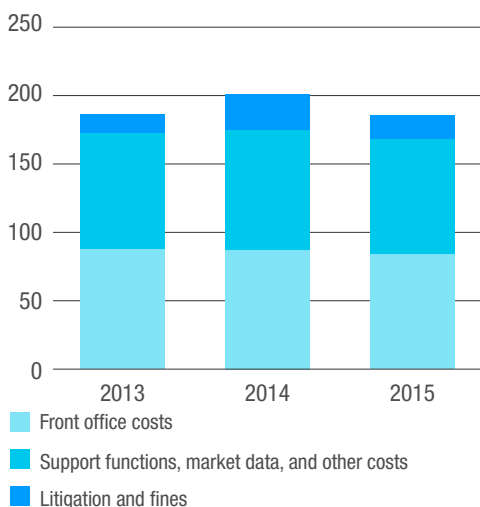
Source: Global capital markets 2016: the value migration, BCG Perspectives, http://on.bcg.com/2tmMjxM

Supporting all this trading activity (and generally reflected in the costs discussed) are many external organizations (and their infrastructure), 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 will typically involve many systems (internal and external to the bank) and parties.

Investment banking, which also includes businesses that are not explicitly related to trading but often result in trades, such as equity issuance, debt issuance, and mergers and acquisitions (M&A), has always been a highly cyclical business, with profits dramatically rising and falling in line with market activity. Historically, one of the key shock absorbers of this volatile business has been staffing costs. Good times meant large bonuses, bad times meant no bonuses and staff cuts. However, since the Great Financial Crisis (GFC) of 2007-2009, trading activity has generally proven to be subdued, normal operating costs have become harder to reduce (in spite of significant reductions in the size of the front office), and control and compliance costs, along with the on-going fines and litigation costs, have grown substantially.[1] Additionally, regulations limiting the banks' trading activities, requirements for additional capital, and restrictive targets for liquidity and leverage have severely reduced the scope for increasing revenues.

It is not surprising that in such an environment investment banks are more focused than ever on reducing operating costs. However, the stickiness of costs has meant that they have had to consider a range of new, and not so new, ideas, such as greater use of off-shoring and outsourcing, mutualization of business functions, digitalization, and the application of distributed ledger technology (DLT).

Given the amount of interest expressed by investment banks in DLT over the past few years it is worthwhile to consider its implications. And, while it is true that many banks, consortia, and fintech companies have undertaken successful trials of the technology, the question remains as to whether it will be able help to reduce the aforementioned costs.

This paper argues that a pragmatic, hybrid approach to applying distributed ledger-like technologies can help reduce the cost of processing trades in financial institutions. However, for that to happen, it does need to be applied in ways that might be quite different to those currently being considered. This is because we believe that even if DLT does take off on a large scale in investment banking it will be part of a mixed environment of centralized and distributed systems for a long period. A pragmatic, hybrid approach would also make it easier for production quality DLT solutions to meet the banking sector's strict requirements in relation to procurement, security, and data privacy.

---

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

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 be capable of integration with a great deal of existing infrastructure. Those who have worked in investment banking IT departments know that, in spite of 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."

A hybrid approach that includes elements of "big data"/analytics, as well as "nudges" of behavioral economics, can help banks tackle one of the major challenges they face implementing large scale change, namely complexity. Complexity makes organizations (and systems) hard to measure, understand, and consequently change. DLT combined with "big data" has the potential to introduce the transparency and feedback loops that are missing in complex organizations. With those in place, there is the potential to "nudge" organizations towards the standardization and behavioral change, which would ultimately reduce costs and operational risk.

2. WHAT DO WE MEAN BY TRADE PROCESSING?

When considering complexity within an investment bank it is very easy to lose sight of what all those systems and departments, spread across so many locations, are ultimately trying to achieve. While it is important not to trivialize any of those functions, it is possible to look at trade processing in a way that makes things much clearer.

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 like, "How much money is being made?", "What risks are being taken?", "What trades could or should we do next?", etc. There are also outputs that are needed to meet the requirements of regulators, accounting standards bodies, and clients.

This does not mean that the optimal infrastructure would consist of simply creating the box labelled "Trade Pocessing Infrastructure" in figure 2 as a single "smart contract."[2] 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 (and hard to maintain) piece of code.

2.1 A model infrastructure

Another key point to understand about the trading infrastructures of most investment banks is that broadly speaking they work. Trades get settled, credit limit utilizations updated, and postings are made to the general ledger. The extent to which that infrastructure is cost effective, flexible, and controlled varies enormously between banks. In spite of the problems, which are to some extent shared across all banks, it is possible to draw a picture of what "good infrastructure" 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 (electronic communication networks), SEFs (swaps execution facilities), and the firm's own external facing trading platforms. This needs to be fast and it needs to be highly resilient. Bursts of tens of thousands of trades may hit this infrastructure over the course of seconds.

Figure 2: Inputs and outputs of trade processing

Inputs: TRADES, ORDERS, REFERENCE DATA, MARKET DATA → TRADE PROCESSING INFRASTRUCTURE → Outputs: Posting to general ledger, Cleared trade, Settlements, Sales credits, Stock position (depot), P&L, Cash position (nostro), Credit risk position, Liquidity position, Trading positions, Market risk position, Risk-weighed assets, Balance sheet, Reported trade, Margin calls
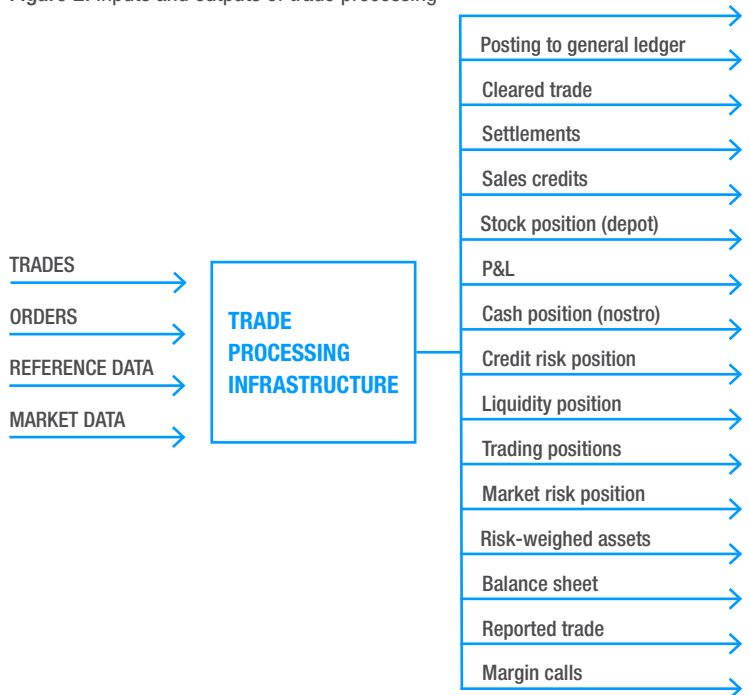
[2] A reasonably commonly accepted definition of a "smart contract" is a code that can run autonomously to enforce and execute the terms of a contract. However, common usage of the term may vary between any programs running in a distributed ledger to genuinely independent objects that could theoretically have their own legal personality.

Those trades need to feed into the bank in as frictionless a manner as possible so that they can be quickly credit checked, executed, hedged, and used to update positions and 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 the even faster algorithmic hedge funds. Resilience and lack of friction is just as important as speed. There is no point 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 sophisticated infrastructure in this area. Highly paid IT professionals have also been hired to build this type of infrastructure in hedge funds, and in some cases they are far more sophisticated than the banks they trade with.

Trades and orders are captured in this layer and are typically fed into a trading system.[3] Historically, the trading system[3] was the place were trades were booked, positions managed, and risk and P&L calculated and viewed, i.e., where traders did their trading. Today, the "trading system" is turning into something of a misnomer with ever greater proportions of trades being executed on external platforms, more trading decisions being taken by automated processes, and risk and P&L being viewed on a cross-asset class basis.

There is still, however, the need for a central view of risk and P&L. There are also the more complex, the less liquid, and the "voice" trades; all of which need people and a trading system. Good trading infrastructure below the low latency layer may not need to work at quite such a frantic pace but it still requires to be well integrated with other infrastructure, to have good quality reference data, and the right business logic. Both internal or 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 assume that trades are then fed from the trading system to the settlement system where trades are then "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 depos and cash in Nostro accounts may be updated

and various exceptions resolved by the 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).

Figure 3: Core trade flow and related systems
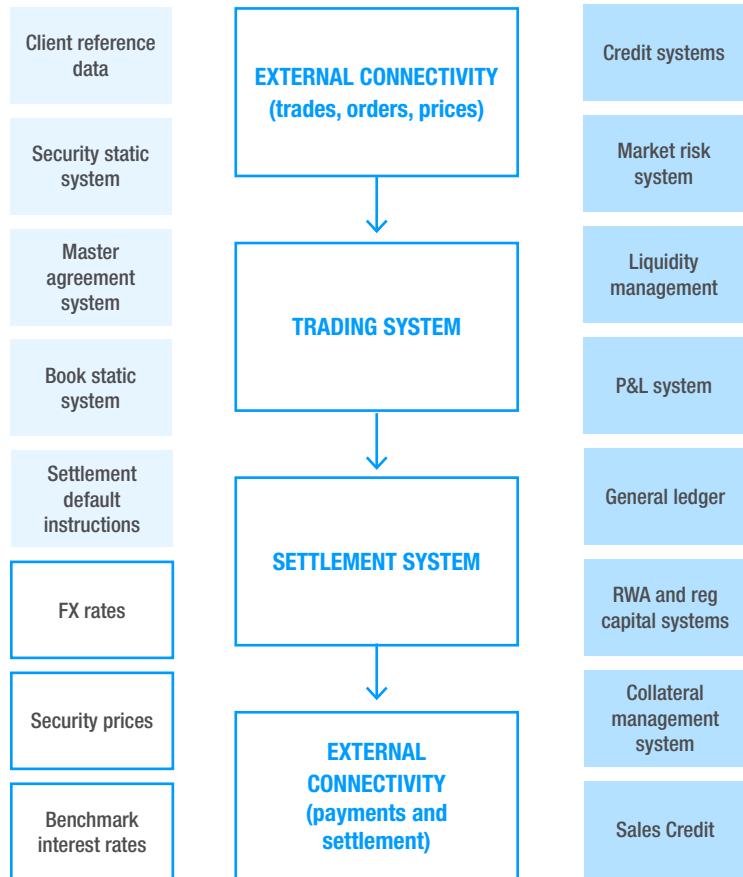
The settlement infrastructure (and often closely related infrastructure for post-trade confirmation and matching) is typically the place where the noise of bad reference data and mis-booked trades becomes apparent. However, if reference data feeds are good, everything upstream is integrated well, and the traders show a disciplined approach to booking their trades, the operational systems can work relatively smoothly. Despite that, counterparties can still inflict pain on the best operations department. For example, they can incorrectly book their side of the trade, they can demand strange quirks in the post trade processing, or simply be unresponsive when errors are found.[4]

---

[3] Trades, quotes, orders, and availability are also generally generated in this layer to feed to external platforms.
[4] Examples include ad hoc changes to netted settlements, requests for non-standard information to be added to confirmation, third party payments of the back of an FX transaction

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 behavior.

Settlement systems share many of the same attributes as the trading systems. There are many mature, high quality systems available from vendors, but even the best of these can be perform poorly due to poor configuration and integration. Some banks build and maintain their own settlement infrastructure, which may reflect the strength of having 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 actually 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 the risk and finance departments. While the core trade processing systems are frequently specific to an asset class, those supporting other functions are generally cross-asset class. P&Ls[5] need to be generated for the bank as a whole, as well as at the level of the trading desk or a book. Risk related to a particular counterparty needs to be viewed across all business lines. Although 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.

Risk, finance, and related systems have exactly 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 GFC.

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.

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.[6] However, a significant degree of data is still exchanged as part of end-of-day batch processes, using a wide variety of methods. There may be SFTP (secure file transfer protocol), 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 lost or mutated.

Other factors that should help the creation of smoothly running infrastructure are standardized messaging formats, such as FpML[7] (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. Adding all these elements together, using consistent market and static data, and making sure that all traded products are fully supported by all the relevant systems, with messages that mean what they say, can create an efficient and low friction infrastructure.

2.2 The problem of trade processing

Investment banks today can process vastly more trades, more quickly, and at lower cost than in previous decades. Products that in previous decades could not even be imagined are traded and processed on a daily basis. Banks would particularly like to reduce their technology costs since they are large in absolute terms and as a proportion of revenues. 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 resources devoted to, frequently fruitless, attempts to improve the infrastructure. Generally, the bigger and more ambitious the project, the greater the risk of failure. Front-to-back re-engineering, "simplification", front-to-back systems, horizontally organized functional systems, and clever (but expensive) middleware programs are all approaches that have been tried multiple times across the banks with very varying degrees of success. There have also been many complete failures. The other

---

[5] Which in a non-banking environment equates to revenues rather than profits
[6] The messaging system uses a built-in data store to persist messages. It does not guarantee that the recipient can successfully process the message.
[7] Financial Products Mark-Up Language

problem is the "noise" that develops 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 staff use the terrible phrase "fire and forget;" send the trade downstream but do not worry about whether anyone can make sense of it or not.
- Lack of reference data systems or failure to connect all relevant systems to those reference data systems.
- Poor trade booking due to the errors of front office staff combined with a failure to encourage better practices.
- Trades and structures that are booked in one systems but not understood by the systems they are fed to (even if the trades are messaged correctly).
- 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.
- The mutation of "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; a theme almost every investment banker would recognize. 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 miss the relevant reference and trade data sourced from golden sources.

"Noise," or "friction," has other indirect costs in addition to the labor costs of paying people to fix problems; it needs more layers of control. In a fragmented, "noisy" infrastructure, those layers of bolt-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.

Regulators, the technical press, and even senior bankers have grown highly critical of the state of investment banking infrastructure. Criticism extending 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' responsiveness to change were the trade reporting requirements under the Dodd-Frank Act (DFA) and the European Market Infrastructure Regulation (EMIR). These requirements were superficially straightforward:

Figure 4: The complexity cycle
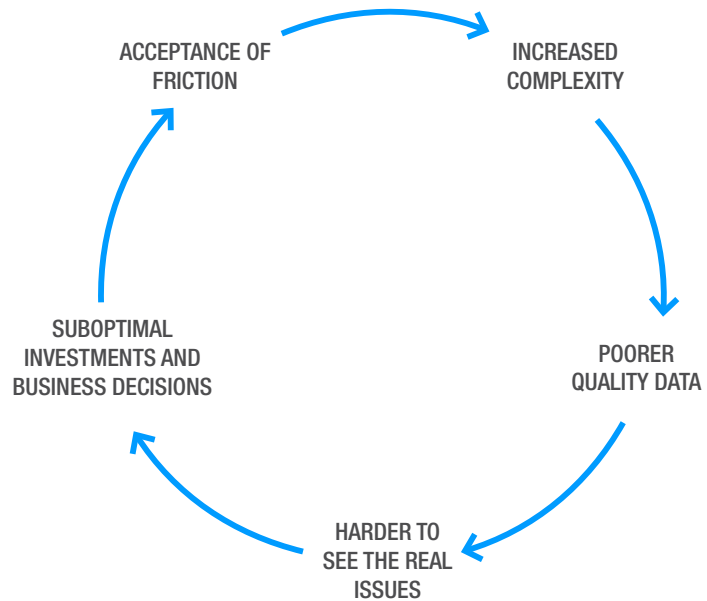
to report both the trade economics of OTC derivatives in near real-time and post-trade data by the end of the day. The exercise proved problematic, if not traumatic, for most banks with costs ranging from the tens to hundreds of millions of dollars per bank.

The complexity discussed 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 wrong for the organization.
- 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 centralization and opportunity for DLT

One of the areas where the greatest progress has been made in trade processing in recent decades has been the efforts to increase standardization and centralization[8] of processing. The huge increases in volumes and variety of products would not have been possible without **centralized** services, such as trading ECNs, the SWIFT network, CSDs, Euroclear and CLS, and **standardization** of product definitions, legal documentation, and message types through the work or organizations such as ISDA, SWIFT, and ISLA.

In recent years, there has been increased desire to extend this approach to the cost challenge. Banks have been more open than ever to the idea of mutualization of their system functionality in areas that they do not consider to be differentiating. However, creating mutualized utilities has not proven easy. Infrastructure in the middle parts of the trade lifecycle have proven particularly hard to mutualize because the systems in those areas have to deal with the most friction-driven complexity.

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 the technical term, complexity is a function of the number of "edges" and "nodes" as shown in Figure 5.

Ideally a technology would be available that:

- Has the same magic ingredient of the best centralized systems, i.e., standardized data models and business logic.

- Deals with the basic drivers of complexity, i.e., reduces the number of edges.

- Supports the mutualization 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, simplifies the settlement process and reduces the time take for settlement cycles.

Figure 5: REPO trade processing (simplified generic architecture)

---
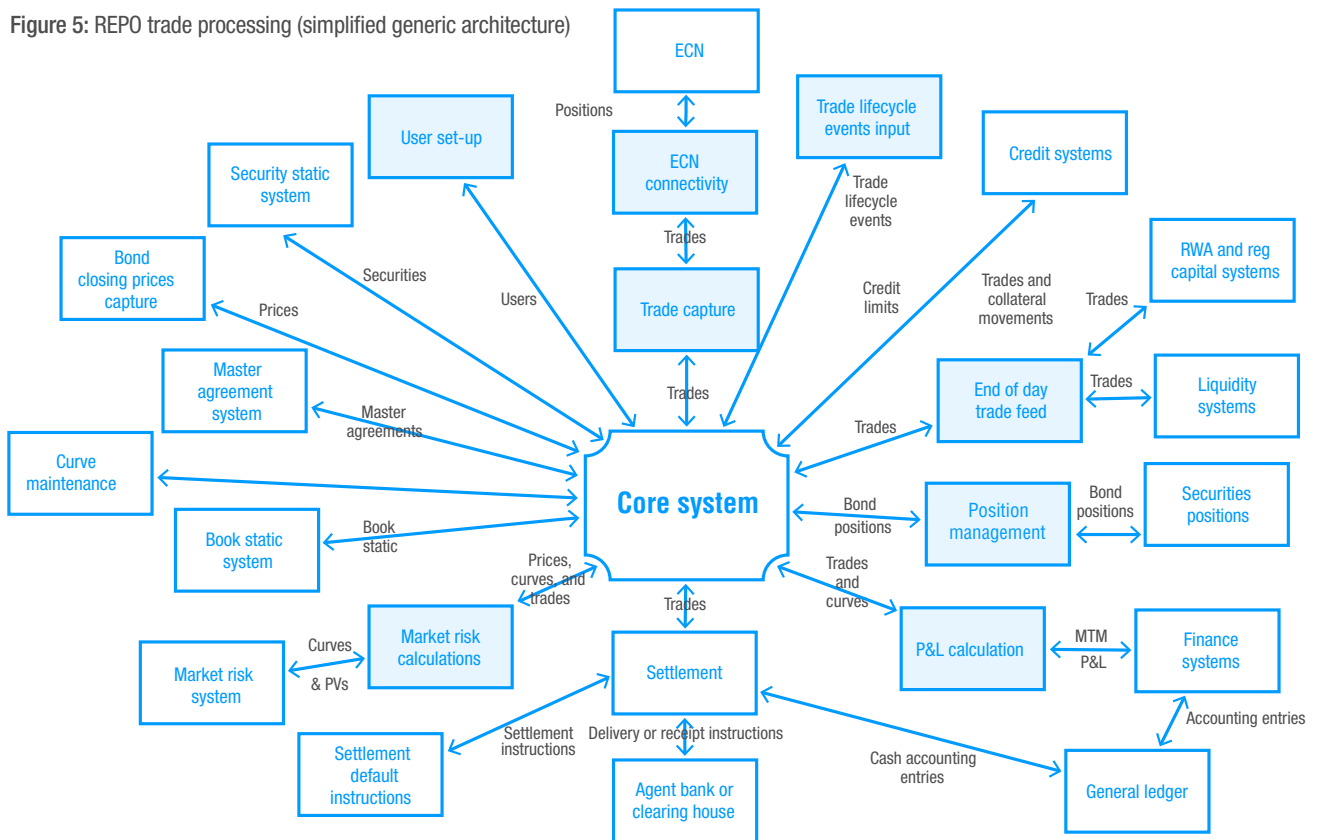
[8] Not something you may have anticipated reading in an article about distributed ledgers

Many people would argue that DLT is that technology. However, the question is how much of that is really 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 that apply consistent business logic to produce the various outputs required to operate a trading business (Figure 2). However, in spite of many proofs of concept in various aspects of capital markets and the major investments made by some DLT-related companies, there are a number of obstacles that need to be worked through to in order facilitate more widespread of adoption of DLT in markets. **None are insurmountable but all could take considerable time and effort.**

A distributed ledger-based solution to trade settlement needs to be tangibly better in terms of cost, control, security, and resilience than financial market infrastructures already in place; much of which works remarkably well, such as the major CSDs and CLS Bank in the FX market.

In ledger nirvana, settlement infrastructure is typically based on the assumption of cash and securities on ledger, i.e., cash or securities legally exist on the ledger or have a "tokenized" representation of assets that enjoy the same degree of legal certainty and settlement finality as the primary record of the assets.[9] 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 performing core processing with all transactions ultimately being recorded on the Bitcoin blockchain. This represents a small 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, a real-world fiat currency needs to be a ledger version of central bank reserves, which require the cooperation of central banks and in some jurisdictions legal changes, or commercial bank created money. Commercial bank money would be the direct economic and legal equivalent of the money that is represented today as positive balances in customers' bank accounts. However, commercial bank issued "ledger money" would still have 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 be needed to make it possible for them to be provided as collateral between banks, to CCPs, and central banks through their repo process, in addition to simply being bought and sold.

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

> "A smart contract that performs the basic mechanics would still need to interact with credit risk, market risk, liquidity management, position viewing, P&L calculation/aggregation, regulatory report, derivative clearing, sales credits, and many other systems."

In many markets, such as spot FX, futures, cash equities, and the more liquid bond issues, the majority of 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 it is the reality of how many markets operate today. There would be great resistance by the markets to any attempt to slow down trading to allow DLT, which is inherently slower, to replace the current pre-trade

---

[9] Tokenization is still largely a concept rather than legally and regulatory verified reality.

Figure 6: From assumptions to solutions

ASSUMPTIONS

- Within and between banks data business logic is already highly distributed, just in an inconsistent way that is prone to the creation of friction in many parts of trade processing
- Complexity makes mutualization or centralization of the existing distributed logic and data, hard, expensive, and risky
- Work towards DLT nirvana will be slow and possibly never reach the destination enthusiasts hope for

THE WAY FORWARD REQUIRES

- A new model for sharing 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 behavioral 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/mechanisms for behavioral 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

infrastructure. A further obstacle is that the post-trade processing costs of electronically executed trades are considerably lower than for the more traditional (and error prone) voice trading. This means that DLT solutions need to be significantly better post-trade than systems and processes that deal with relatively trouble free "e-trading."

Ledger nirvana would also make many proposed distributed ledger-based systems fall within the scope of the Bank for International Settlements' (BIS) "Principles for financial market infrastructures."[10] These principles are incorporated in law in most jurisdictions and are justifiably demanding. They represent a high 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 DL does not automatically remove the need for a general ledger, the use of cryptographic techniques does not make a system more secure from a bank's perspective.

In ledger nirvana, the trade is the settlement. A trade is booked and value exchanged. However, this creates significant problems for today's business models, which cannot simply be wished away by the DLT enthusiast. Most of capital markets works 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 marker 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 purchasing. 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 need to improve trade processing, as well as reducing 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?

Figure 6 presents the visual construct of what happens if you accept the assumptions implicit in the analysis above.

The model described below incorporates elements of DLT technology, data analytics, and existing software tools to meet the problems described at a technical and organizational level.

---

[10] http://www.bis.org/cpmi/publ/d101a.pdf

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" (PDA) 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 and validate them, as well as either create a new object (in accordance with the PDA) 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 will also 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 PDA

The PDA provides the basic set of rules parties need to accept regarding the processing of trades they agree on, on a collaborative, distributed basis.

Figure 7: Sample product definition agreement

PRODUCT DEFINITION AGREEMENT	
Product Name	Spot FX
Product Code	FXSpot
Data Model	ISDA FXSpot
Party A Signed Date/Time	22/11/2016
Party B Signed Date/Time	19/11/2016

ATTRIBUTE OWNERSHIP	
ATTRIBUTE	ATTRIBUTE OWNERSHIP
Trade Date	Consensus
Settlement Date	Consensus
Currency Pair	Consensus
First Currency Amount	Consensus
Second Currency Amount	Consensus
First Currency Settlement Instruction	Party A
Second Currency Settlement Instruction	Party B
Valuation	3rd Party

FUNCTION OWNERSHIP						
FUNCTION	PARTY A -APP/SERVICE PROVIDER	PARTY A -APP VERSION	PARTY A -OBJECT	PARTY B -APP/SERVICE PROVIDER	PARTY B -APP VERSION	PARTY B -OBJECT
Valuation	Open Gamma	1.02	Trade	Open Gamma	1.02	Trade
UTI Exeception Resolution	Lebowski Inc	2.1	Trade	Lebowski Inc	2.1	Trade
Matching	Calypso	14.1	Trade	Calypso	14.1	Trade
Settlement Engine	Calypso	14.1	Trade	Calypso	14.1	Trade
Netting	Calypso	14.1	Settlement	Calypso	14.1	Settlement
Collateral Management	Cloud Margin	3.2	Trade	Cloud Margin	3.2	Trade
Payment Instruction	Bank Internal	4.5	Settlement	Calypso	2.3	Settlement
Operational Risk Monitor	E & Y	1	Trade	KPMG	9.7	Trade

A consensus mechanism in DLT, such as the "notary" in R3's Corda allows all relevant parties to "sign" that they are agreed to the processing rules and data structures.

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 e.g., one party may "own" the population of their own settlement instructions, which are then accepted by all relevant parties for use in the settlement process. For others, such as a valuation, a third party may have the obligation to populate the data. Finally, many attributes, such as trade economics, will be owned by both trading counterparties and the attribute is only recognized as correct when they match.
- Functional ownership records which elements of the trade lifecycle will be carried out by a specific system or party. Both parties may agree to use a specific smart contract, or a cloud version of existing vendor system, a centralized service, or they may agree to continue using their systems (accepting increased risk of differences).

3.3 The golden container

Trade and event data is stored in the "golden container." There has been a seeming endless quest by many in banks for a "golde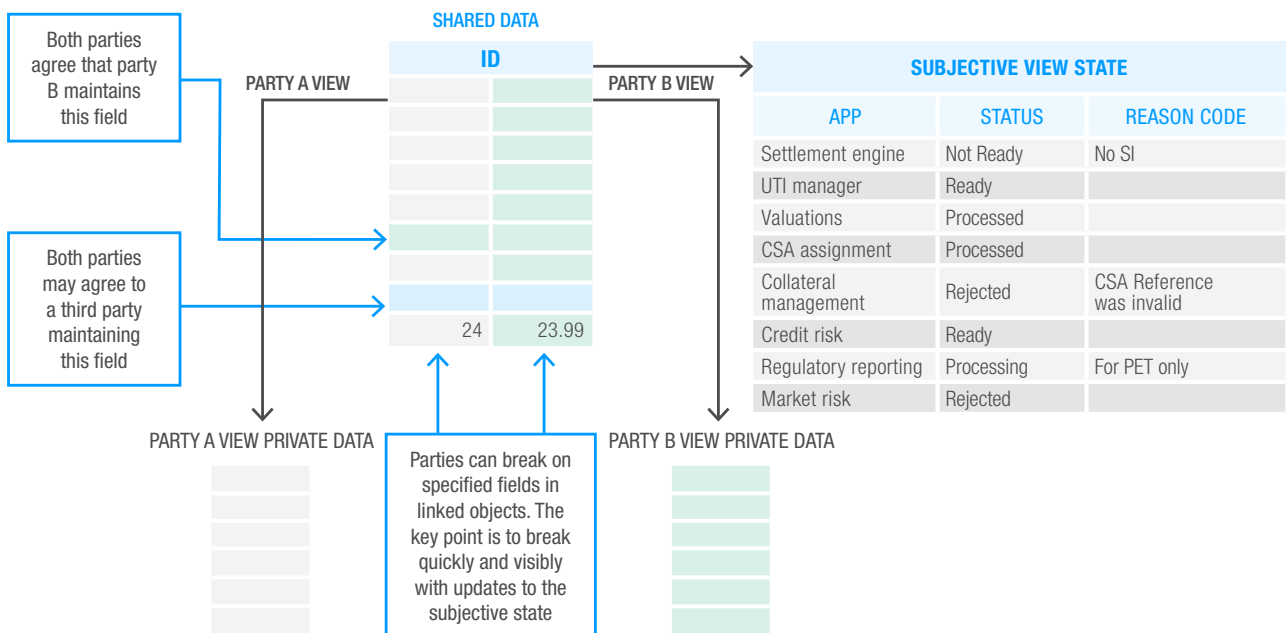n trade" record. The general idea is to create a single version of the truth for a trade (either at an organizational or a market level). Unfortunately, the idea tends to break down unless it involves a centralized system to not only store the golden record, but also carry out all of the processing. 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 market's proof of concepts that assume either consistency of the trade from the outset or that the trade is either in an agreed or not agreed state. It would still 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 the 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 to create 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 it to their counterparty (or any other agreed relevant party).

Figure 8: The golden container

| SUBJECTIVE VIEW STATE		
APP	STATUS	REASON CODE
Settlement engine	Not Ready	No SI
UTI manager	Ready	
Valuations	Processed	
CSA assignment	Processed	
Collateral management	Rejected	CSA Reference was invalid
Credit risk	Ready	
Regulatory reporting	Processing	For PET only
Market risk	Rejected	

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

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 to avoid mutation through the addition of data superfluous to one party or the "overloading" of fields.
- **Subjective view states** make it possible 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 (i.e., the front office, which requires only a calculated PV), but may not be valid from another function's perspective (i.e., the collateral management system may require the relevant master agreement to be assigned to the trade).

**Figure 9:** Relationship between key objects
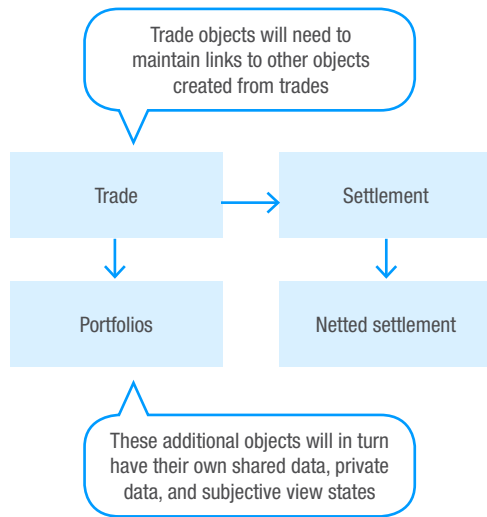
3.4 Data capture

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

The first key step to adopting a more distributed approach to trade processing is to recognize 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 are already agreed on the

Figure 10: Data Capture
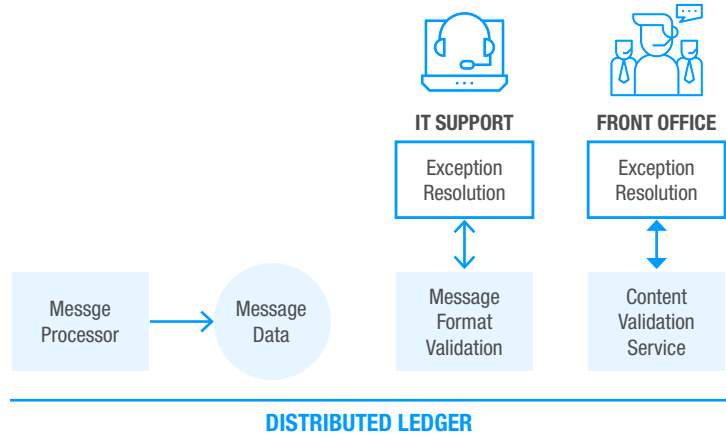
trade economics, such as an ECN, or from an existing trading system where trades are input manually (and unilaterally) by one party's front office.

The data capture mechanism needs to be able to process any popular message format used today. To maximize flexibility, trading counterparties should be able to agree to use a third-party message validating/translation service (see PDA). There are several proven products on the market that 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.[11]

Figure 11: Data distribution

---
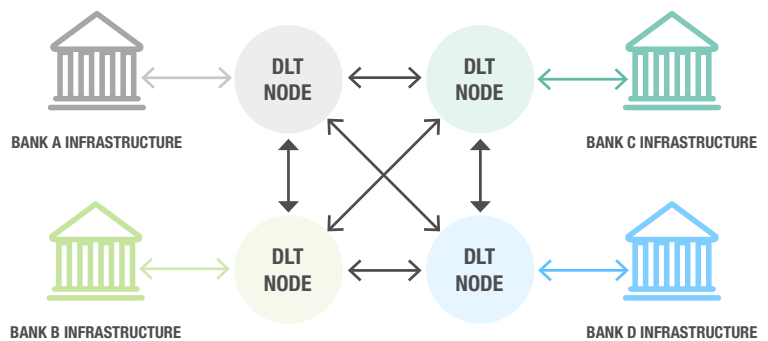
[11] This includes Xceptor's DataHub or Broadridge's Message Automation tools

3.5 Distribution

Nodes[12] 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.

3.6 The status monitor

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

Some banks have attempted to create similar tools to benefit from potential control and operational efficiencies, but they are generally held back by the fragmented state of their architecture.

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

Figure 12: Examples of status monitor

| PORTFOLIO VIEW					
FUNCTION	NOT PROCESSED	SENT FOR PROCESSING	PROCESSED	ERROR	NOT APPLICABLE
Trade Capture	0	234	234	0	0
Trade Validation (Technical)	0	232	231	1	0
Trade Validation (Business)	0	123	122	1	0
Matching	2	80	74	6	154
UTI Exception Resolution	31	123	123	0	200
Settlement Engine	0	143	142	1	0
Valuation	1	154	123	3	43
Trade Reporting	3	120	118	2	114
Clearing	2	32	31	1	202
Netting Engine	0	123	122	1	32
SWIFT Settlement Gateway	1	13	23	12	123

Provides a filterable view of status. This is likely to be used by operations, control functions, and IT

TRADE VIEW		
FUNCTION	STATE	ERROR TYPE
Trade Capture	Processed	
Trade Validation (Technical)	Processed	
Trade Validation (Business)	Processed	
Matching	Processed	
UTI Exception Resolution	Not Processed	
Settlement Engine	Error	Missing Settlement Instruction
Valuation	Sent for Processing	
Trade Reporting	Processed	
Clearing	Not Applicable	
Netting Engine	Not Processed	
SWIFT Settlement Gateway	Not Processed	

The ability to drill down to trade level means anyone involved in any part of processing the trade can make a judgment on the overall operational risk. A trade level view could also display information on trade economics, etc.

[12] A system that manages the communication of messages between all parties sharing the same distributed ledger based system.

Figure 13: The overall ecosystem
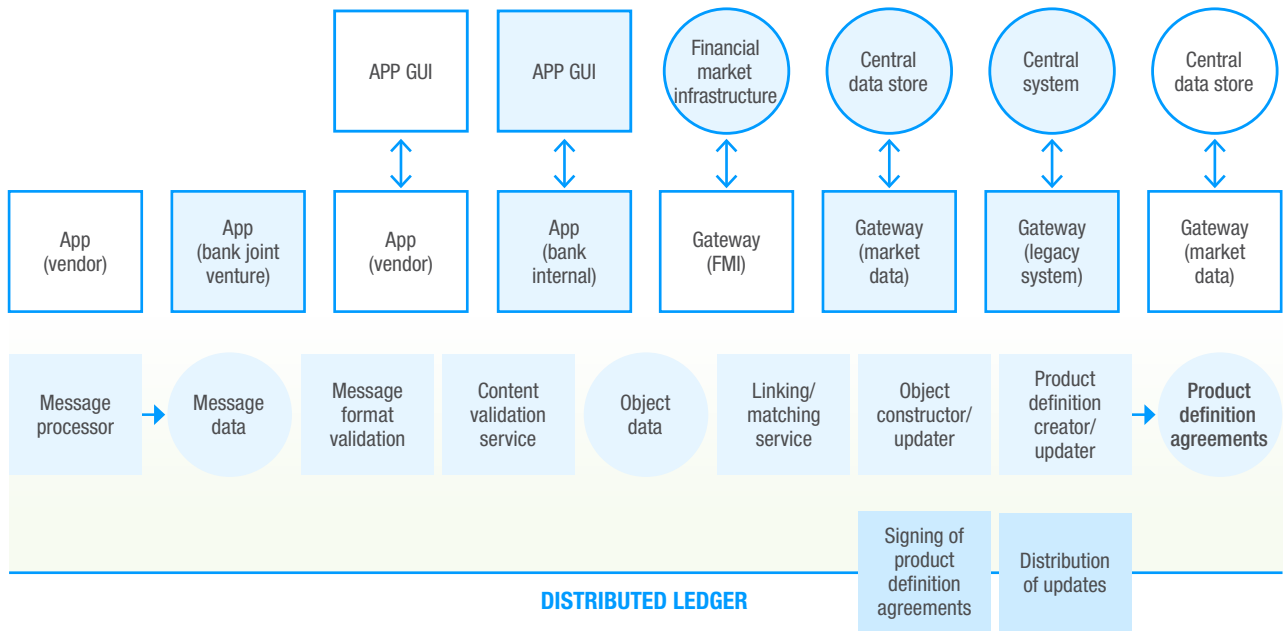
3.7 Functional model

Overall, this design depends largely on the core capabilities of DLT systems and the ETL tools already in use in many banks. The new functionality consists of a set of small and relatively simple components.The "object constructor/updater" determines whether a new object, i.e., 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 from a counterparty's (or other third party) data. If it finds a sufficient degree of consensus it will merge the objects into a single golden container (see below) that have both parties' versions of the object in a single data structure. There will also be the option for both parties to "force match" two objects into a single golden container, if they both agree their versions represent the same trade, even if they potentially disagree about some of the details.

A "service notifier" will inform all relevant connected applications 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, such as settlement or trade reporting. The service notifier will also inform the relevant services

if an updated object has sufficiently changed 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.

3.8 The feedback loops – cost and control

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 PDA) about the "meta-workflow" of trades, such as when it was matched, when it was complete enough for settlement, and the time delays in processing

Figure 14: Integrating analytics components
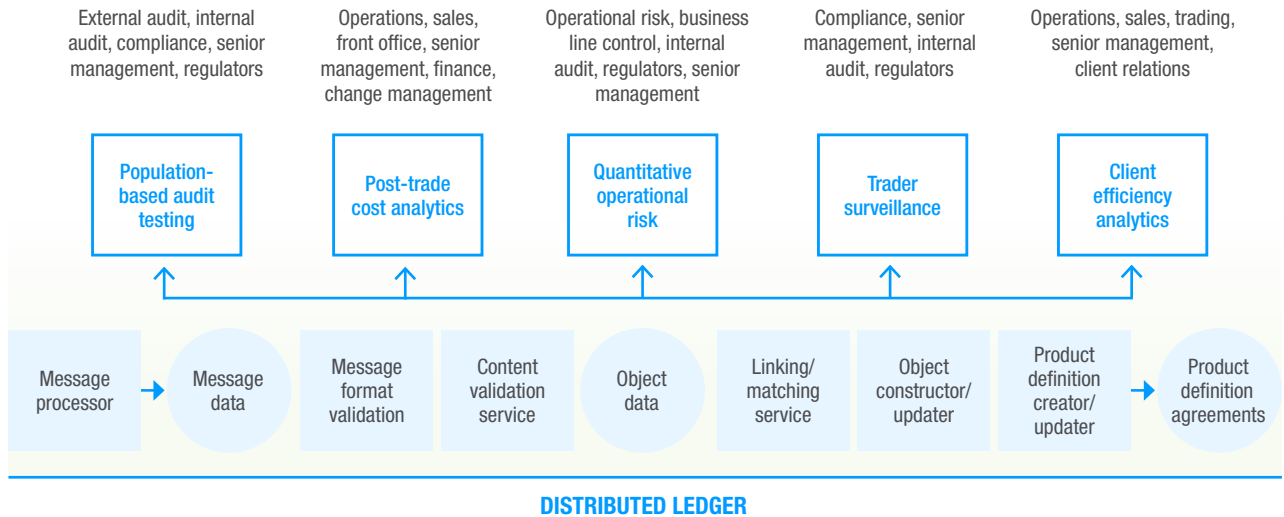
The tools, which could be produced in partnerships with organizations specializing in control and/or analytics, such as audit firms, include:

Population-based audit testing: currently the external audit process is based on samples of trades and does not trust the bank's own internal records. If there is a trustworthy source of shared trade data, external auditors could test the validity of all trades recorded by one party rather than just using a sample, i.e., population based testing. Potentially, tools could be deployed that do continuous real time auditing of trade data.

Post-trade cost analytics: the main drivers of the human costs in trade processing are exceptions to straight-through processing (STP), 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 estimateS 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.

Trader surveillance: current systems used to detect rogue traders typically look for patterns of unusual behavior in trading and trade processing. An approach

"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."

based on the use of the status monitor allows consistent business logic to be used across multiple 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 each client, 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 an explanation of the mechanisms causing problems in trade lifecycle processing and the techniques that could be used for dealing with them. Many problems are clearly not technical but socio-economic or political in nature within banks. Nobody sets out to design bad infrastructure, but a succession of decisions, which are optimal for one specific

function or business lines, can progressively create complexity and drive entropy in the infrastructure and the organization. Complexity reduces understanding of the systems, which drives more sub-optimal decision-making, in turn 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 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, 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 that 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 undisruptive way. It also provides the feedback loop that can objectively measure the success of different approaches, whether they use DLT, or current technologies, or simply involve organizational or process change.

ABOUT CAPCO

Capco is a global technology and management consultancy dedicated to the financial services industry. Our professionals combine innovative thinking with unrivalled industry knowledge to offer our clients consulting expertise, complex technology and package integration, transformation delivery, and managed services, to move their organizations forward. Through our collaborative and efficient approach, we help our clients successfully innovate, increase revenue, manage risk and regulatory change, reduce costs, and enhance controls. We specialize primarily in banking, capital markets, wealth and investment management, and finance, risk & compliance. We also have an energy consulting practice. We serve our clients from offices in leading financial centers across the Americas, Europe, and Asia Pacific.

To learn more, visit our web site at **www.capco.com, or follow us on Twitter, Facebook, YouTube, LinkedIn and Xing.**

WORLDWIDE OFFICES

Bangalore	Hong Kong	Singapore
Bratislava	Houston	Stockholm
Brussels	Kuala Lumpur	Toronto
Chicago	London	Vienna
Dallas	New York	Warsaw
Dusseldorf	Orlando	Washington, DC
Edinburgh	Paris	Zurich
Frankfurt	Pune	
Geneva	São Paulo	

CAPCO.COM 🐦 f ▶ in ✗

CAPCO