

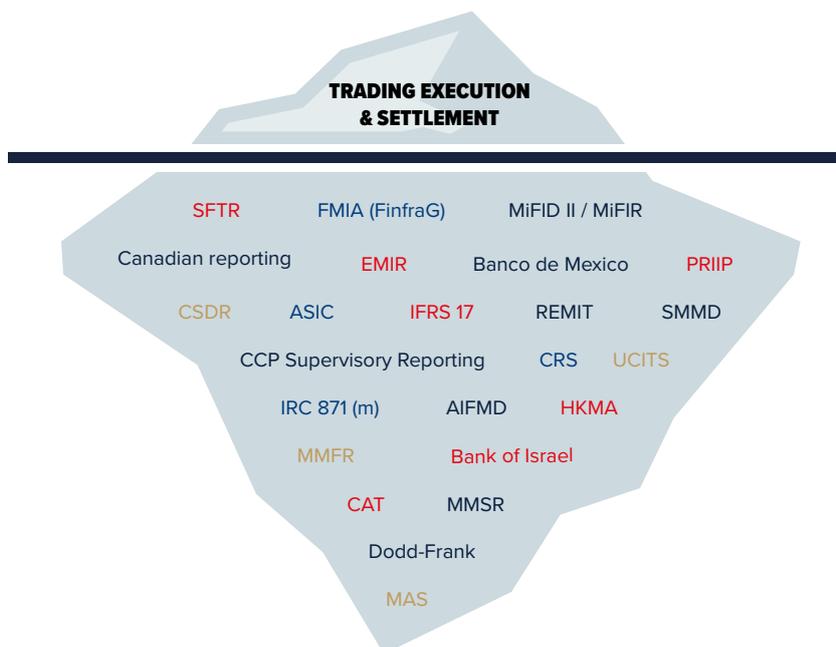
THE CHALLENGE OF REGULATORY TRANSACTION REPORTING

When it comes to regulatory reporting, there are many different challenges, several of which will be discussed here. To be clear, this article talks about transaction reporting, specifically post-trade regulatory reporting regimes, rather than trade reporting. They're related but they pose different challenges.

By Fiona Hamilton, Research Director, Volante Technologies

MULTI-ASSET, MULTI-GEOGRAPHICAL, MULTI-REGULATION = MULTI-COMPLEX

Since 2008, financial institutions have experienced an avalanche of regulation. The US introduced the Dodd-Frank act which was followed by EMIR and MiFID II in Europe. Of course, as organizations are now invariably global, it isn't just US or European regulations that need to be tracked and adhered to. Tracking regulations in Hong Kong, Singapore, Canada and Australia also becomes a necessity for many global firms. Add to this the complexity of extra-territoriality with regulation coming out of the US that affects everyone trading in US stocks and fixed-income instruments, and similarly EU regulations which affect the US, Australia, Singapore and beyond.



THE PASSAGE OF TIME

Historically, regulation takes time to pass through the appropriate legal processes before being introduced into law, creating a time lag from when the need for regulation is born and when it is introduced. If we look at this last decade since the 2008 financial crisis by way of example, we are only now seeing a big wave of regulation coming into force. MiFID II, the successor to the 2004 MiFID, now includes MiFIR which directly addresses regulatory reporting and only came into force in 2018. Now, following fairly quickly behind is SFTR in the securities finance area, which will affect a huge amount of capital markets transactions in repos and stock lending; MMSR is to be implemented for money market statistical reporting and CSDR for central securities depositories. All will have an enormous impact as they cross multiple asset classes: equities, fixed income, both exchange-traded and OTC derivatives and foreign exchange swaps and options.

THE CHALLENGE OF LEGALESE

Legal documentation also compounds the challenge. Legalese is not written by bankers or other experts in financial instruments, trading or post-trade transactions – it is written by lawyers, in legalese, as it must be, because it comes from legislators. But legislative terms are often problematic to translate into plain English. For instance, it is often difficult to define a specific instrument type or a particular attribute of an instrument. An example of this is the term 'maturity date,' which, for many instruments, is a meaningless term. How to interpret the legal documentation and implement something which is going to be compliant therefore presents itself as the next challenge.

THE CHALLENGE OF COMPLYING WITH THE TECHNICAL IMPLEMENTATION

Typically, we see different levels in terms of regulatory rollout: level one documents, level two documents and then lastly, level three – the regulatory technical specifications (RTS). The RTS is the definitive artefact which defines from a technical perspective, which message standard and orchestration is required to exchange - crucial - information between a financial institution and a regulator or an intermediary. The intermediary could be an ARM or it could be a trade repository, but the RTS will indicate the minimum data elements which need to be reported. RTSs can often be produced and agreed just prior to the regulation legally coming into force and unfortunately, sometimes, afterwards.

For major regulatory initiatives such as MiFID II and Dodd-Frank, banks will typically start to assess how the regulation will impact their business several years before having to be compliant. There will clearly be an effect on the business but equally important, they need to understand what the effect will be from a technical standpoint. If the RTS is published three or four months before a regulation becomes law, then it is impossible for a bank to apply those RTS specifications to the standards within that timeframe. With no clear direction on how to technically implement the advice they will have received from a myriad of advisers, banks are left with the challenge of significant change that must be implemented within a short period of time.

PRIORITIES OF BUSINESS VS TECHNICAL IMPACT

There has been a lot of industry discussion around the unbundling of research from fees, the need for greater transparency and so forth. This has resulted in the more technical areas such as transaction reporting being left until last because the devil is in the detail [aka small print]. Mandating that research can no longer be bundled and a fee must be charged has a clearly defined business impact and business managers, legal counsels and or consultants will have eagerly assessed the impact on business models. However, looking further down the food chain, at transaction reporting to be precise, and the approach has been to leave the technical elements until last.

GREY AREAS IN REGULATION

As we moved closer to the implementation of MiFID II, there were many apparent grey areas with one national competent authority stating, “We will not be providing any further clarification because if we did it would confuse people.” Leaving the implementors - banks and vendors - with unanswered questions around how to create files to report via the stipulated mechanism being but one of many. It's fair to say that we as market practitioners are aligned with the Regulators' desire for a level playing field, but there is undoubtedly a disconnect between technical, legal and business processes within the regulatory ecosystem that need to be addressed.

From participation in numerous working groups and industry forums, it is easy to get the impression that for the initial period of any regulatory implementation a 'fudge factor' is positively engendered. This is the period in which things are largely unknown and wholly complicated, with nobody - unsurprisingly - willing to stand up in public and say - taking into account the multiple market practitioners within our capital markets universe - we are not 100% certain we've got it right. We will therefore make the regulation loose enough to allow for some 'wobble room'. If it were the case, we might end up in a place where one interpretation of something might be X and another interpretation might be Y, but that could be deemed acceptable even if those two interpretations were mutually exclusive.

THE CHALLENGE OF REGULATORY TRANSACTION REPORTING

REPORTING AREA	KEY DATA ITEMS	REPORTING PARTY	TO WHOM	VIA	WHEN/HOW OFTEN
Pre-Trade Transparency (20,21)	Big and offer prices	Trading venues (MTF, OTF, SEF), investment firms, systemic internalizers, (no delegation)	Public	Trading venue (CTP)	Continuous
Post-Trade Transparency (20,21)	Private, volume, time of trade	Trading venues, investment firms, (no delegation)	Public	Approved Publication Arrangement (APA)	Real time (ASATP)
Transaction Reporting (26)	Trade data sufficiently detailed to allow NCAs to monitor for market abuse, regulatory breaches and systemic issues	MiFID firms, trading venues (for non-MiFID firms) (can delegate)	Competent Authority	Approved Reporting Mechanism (ARM)	T+1

As we move through any regulatory cycle, the landscape typically becomes gradually more circumscribed. Looking back for instance at the first EMIR-related fine, it only came in 2017 - several years since the regulation was first enforced and post the wiggle-room period. History has since repeated itself and we have seen the same with other regulations and the ensuing fines.

The regulators are right of course and there is a clear case for a settling-down period due to the complex nature of our capital markets. Conversely however, it comes down to timing; the wiggle-room can be a good thing – initially – but the danger is that it might delay the process of putting in place proper regulatory reporting solutions as well as other aspects of regulatory compliance.

COBBLED TOGETHER SOLUTIONS

Banks and financial institutions work on shareholder value. If US-based you will need to report to your shareholders on a quarterly basis, If UK-based, it might be once every six months. But in either case, reporting to shareholders will always take priority and consequently, implementing support for regulatory reporting may appear, from an outsider's point of view, to become side-lined almost to the point of only doing the minimum now and worrying about fines further down the line. The tacit understanding that you're not going to be fined in the first couple of years compounds this situation. After all, why spend anything other than the barest minimum to achieve compliance but not complete compliance?

As a result, we are seeing 'cobbled together' solutions built in-house by banks and within specific silos; an equities floor will have its own solution under Dodd-Frank and MiFIR, as will fixed income, OTC derivatives, exchange-traded derivatives and FX. The attitude has tended to be to explore how much can be delegated or outsourced to a trade repository or the ARM so they become responsible for compliance. This approach however will be increasingly harder and harder to pull off because several regulations emanating from the EU now dictate that firms take full responsibility for their own reporting mechanisms and some initiatives like CSDR, will stipulate the need for certain data formats from end-to-end, meaning firms will be unable to use existing proprietary data formats. This will cause a huge challenge for some of the trade repositories and the ARMs (not to mention the firms). Whereas the European Union is trying to push all parties, from the absolute originator to the national competent authority, to use ISO 20022/15022 compliant data the whole way through, our industry is still some way from achieving a one-standard-fits-all approach.

MOVING BEYOND PROPRIETARY FORMATS

Looking at the issue of proprietary formats further, it is not unlikely that a firm has had to make use of developers

from other parts of their business in order to build a solution based on a proprietary data structure – which is, arguably, just about good enough to pass on to the reporting mechanism. However, the realisation then sets in that the code that was used needs to be changed or that clarification is required on what needs to be reported but by then, none of those programmers either work within the organization or are unavailable due to other internal commitments. And here's where the challenges of maintaining a reporting framework really come into play. The technical response will almost certainly involve the use of the dreaded word 're-factoring' which is technical jargon for when a programmer/engineer explores code and finds they are unable to understand how or why it was written or indeed, how it works. The de-facto solution will be to re-write the code. And, beware – 're-factoring' means re-writing, bringing with it all the testing overheads, not to mention development, risk and associated costs entailed.

Regulation sometimes does not imply a vector of more reporting; for example, due to changes in administration, in the US, the Dodd-Frank regulatory framework is changing and the logic for data that needs to be reported and thresholds are moving. The issue with a proprietary implementation is that the logic is hard-coded which will cause massive maintenance issues as regulations evolve. But also interesting to observe is that they can also evolve backwards as well as become more encompassing.

ANTICIPATING FUTURE REGULATORY CHANGES

The secret is to anticipate that we're not at the end of regulatory cycles; there will be more. There will be more clarification, and some could go backwards and then maybe four years later will come forwards again. Granted that "forwards" and "backwards" are purely subjective terms depending on your opinion of regulation. However, the fact remains that regulation must be embraced, it is here stay, it's dynamic. It can cost you a considerable amount of money if you fail to comply so it is a false economy to act on it later. Frankly, it's debatable whether you're doing your fiduciary duty if you don't fix something now if you know that the problem will re-surface further down the line and, result in a fine. Historically some people involved in implementing regulations internally within firms were making decisions to do things in a more cost efficient (aka cheap) way knowing that they will need to be fixed later, in the knowledge that they were probably not going to be responsible or deal with the outcome. Changes in other regulatory frameworks such as the UK's FCA's Senior Managers & Certification Regime (SM&CR) are however increasingly focused on the culpability of management that might indulge in that thinking. Senior managers and executives will increasingly be liable to be personally fined or sent to prison if it's been proven that they wilfully decided to skimp on a budget. A fact that should concentrate minds.

REGULATORY COMPLIANCE - SOLUTIONS

AI AND MACHINE LEARNING

Those are the challenges. Are there solutions? The short answer is, none that are easy. Intrinsicly, the challenges are hard but arguably when it comes transaction processing it's just orders of magnitude harder. Solutions must be pragmatic by both embracing the challenges laid out above, but also the opportunities that they can uncover.

There are certainly some forward looking initiatives out there. For example, the FCA invited interested parties to participate in a TechSprint to investigate the feasibility of Model Driven Machine Executable Regulatory Reporting. Certainly, from a purist perspective this makes a lot of sense. But while a worthy objective, the devil is always in the detail. The likelihood of overcoming the practicalities of writing regulation which can automatically generate the code to execute the reporting or the compliance, is still some way off. Yes, it can be done, however if you are a global bank looking at millions of pages of regulatory text implementing hundreds of thousands

THE CHALLENGE OF REGULATORY TRANSACTION REPORTING

of rules across many asset classes, how likely is it that any time soon all of that text will be re-written?

The challenge is that while legal documents are obviously very tightly worded from a legal perspective, they often do not in the world of regulation use the same terms that financial firms use when operating the execution and settlement of transactions. Essentially, a waterfall approach is taken to publishing regulations, and not many legal people understand the terminology that a technical business analyst would use and vice versa - factors which are essential when implementing compliant solutions.

For the automation of regulatory reporting there needs to be clarity, an exact equivalence between legal terms and the terminology that a business analyst would use to communicate with a technical implementor/programmer. Terms like 'rate,' 'date,' 'market' etc. are all very well but to be executable, each term must relate to a specific data element. What sort of date, what sort of rate, etc. It is anticipated that the application of artificial intelligence and machine learning will play a part but as things stand, it's not going to be possible to automatically generate code any time soon across all regulations. Perhaps the next generation of regulatory technology may begin to be able to do this, but even then, to achieve that, the legal profession would have to work with the business analysts to craft the language that would form that bridge between legalese and market practitioner. Whereas arguably a sensible approach, it could also conversely result in as many business analysts on the team as lawyers. It is great way forward and should be fully endorsed, but in the interim, for the next five to ten years, which is the usual cycle of technical replenishment, is there a way we can achieve a halfway house?

DEVELOPING SCHEMAS FOR ELIGIBILITY CRITERIA

It should be possible to take the legal documents that define the regulations and develop and publish what might be termed 'schemas for eligibility criteria'. Ideally these would be open-source to the community at large and would be either regulatorily specific or potentially able to enable reporting under more than one regulatory regime for the same instruments. So if, for instance, one leg is in the USA and the other leg is in the EU, they could be reported under both MiFIR or EMIR and, or, Dodd-Frank.

Ideally, the schema would define the data elements which constitute the eligibility criteria for any regulation, such as Buy/Sell, Domicile, Entity type, Dealing Capacity, Trading Venue, etc. which could be represented by Yes/No flags. Each entry would then need to be defined further, under which regulation(s) / flags a transaction would then need to be reported under. Lastly, each entry would then need to specify the mandatory reportable fields. This approach would work especially well for reports that need to be delivered in ISO 20022 format given that there is significant overlap in the elements that need to be reported. This approach also caters for the instance where reporting is via a proprietary standard as there could be separate schema for different asset classes where ISO 20022 does not dominate, such as reporting to Trade Repositories. What is key is that they capture the key criteria involved with evaluating, what, where, why and who to derive what to report and in what format.

However, what happens when a firm has to report fields one, three, seven, sixteen, twenty-four, thirty-four... what is the mandatory reporting format? Is it to a trade repository or in a CSD format? Is it reportable under an ISO 20022 format? You're not defining the actual message format of what has to be reported, but what has been created and populated by experts in the field which can then be referred to by the technical teams that have to create the final reporting messages. The diagram below shows this type of approach. For simplicity it is in the format of a spreadsheet snapshot but in reality the layout would be defined by an XML Schema and the actual criteria held in an XML format file which complies to that schema.

Conditions for Scenario													Fields to be reported							
Report Status	Buy/Sell	Domicile	Entity	MIFID II investment firm	Conditions for transmission specified in Article 4 satisfied?	Capacity/Condition	Transaction Type	Reporting Obligation for Instrument	Trading Venue	Instrument Type	MIFID II reportable	EMIR reportable	REMIT reportable	"MIFIR Field 1"	"MIFIR Field 2"	"MIFIR Field 3"	"MIFIR Field 4"	"MIFIR Field 5"	"MIFIR MIFIR Field 6"	"MIFIR Field 7"
New	Buy	Firm is dealing within the EEA	Firm is an Investment Firm	Yes	No	Dealing on own account	Transaction is a purchase (buy/sell) of a financial instrument	The instrument has an under-lyer that is a financial instrument traded on a trading venue	Within the EEA and is a Regulated Market	Transferable Security - Listed Derivative	Yes	Yes	No	Report_status	Executing_entity_identification_code	Type of ID of the Beneficiary	Trading_date_time	Trading_capacity	Quantity	Notional
New	Buy	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	New	Yes	LEI	Yes	DEAL	Yes	No

OUTSOURCING AND COMMERCIALISING ELIGIBILITY CRITERIA

This could be an open source eligibility criteria schema that could be populated like reference data. Firms pay for SWIFTref, Acuity, DTCC Avox, Reuters, Bloomberg... firms pay for data, so why could the same model not apply here? Firms spend huge amounts of time and money analysing regulatory texts so why not mutualize the costs of at least the first part of the implementation process?

There are many consultancies both large and niche that have expertise in specific areas - exchange-traded derivatives in one particular market, or credit default swaps or equities in another market or across markets. Why can't eligibility criteria be a commercialised offering where expertise is monetised in that middle layer? That would then allow firms or vendors to build solutions which look-up the eligibility criteria, perform the code-generation and the visibility and the management reporting and the actual reporting. Then it's dynamic, and if/when a regulator releases new information, that reference feed of eligibility criteria could be easily changed.

To achieve truly model-driven code-generation, eligibility criteria should be dynamically looked up and the output changeable in order to achieve compliance. It would not be fully automated between the legislative wording and the eligibility criteria, but automated from the eligibility criteria to the actual reporting. Thus, it becomes dynamic in the same way of changes to an ISIN or for any other reference information such as LEIs. By adopting this architecture, it is easy to see how commercial services and microservices architecture could move eligibility into the API economy.

IMPLEMENTATION OF ISO 20022 BASED CANONICAL MODELS

Outsourcing the eligibility criteria can potentially be viewed as one leg of a solution. The second leg of an extensible and maintainable solution is the implementation of an ISO 20022 based canonical model. Certainly, for most markets and regulatory regimes ISO 20022 is increasingly becoming the dominant standard, but different regulators and markets will implement it slightly differently. So, for technology and products to be able to process that eligibility criteria, the key is for them to hold a repository of the formats that need to be communicated and the orchestration or the choreography between sending data in the appropriate format and receiving it back. You might have 16 different messages to transmit to the FCA than you would have to the SEC or the German regulator or HKMA, but increasingly they are likely to converge on ISO 20022. However, many regulations are based on elements of FpML as a standard, or FIX, or proprietary CSV, but by and large the trajectory does seem to be ISO 20022. Whatever your choice of standard, all three of them are extensible, so the restrictions on the code words which are allowed in fields, the enumerations and cross field rules, can be altered. A canonical model based on ISO 20022 takes this into account by being extended if necessary to support data elements not supported in the published ISO 20022 data dictionary.

BENEFITS OF A MODEL DRIVEN APPROACH

The main benefit of this approach is a common “lingua Franca” regardless of asset class. In other words, the combination of the eligibility criteria and the extensible canonical model facilitates the final leg of the solution approach which is model driven code generation whereby the criteria models and the canonical model are brought together with models of the actual reporting formats with graphically defined validation and mapping rules to automatically generate the output code in whatever technical architecture is required. Perhaps Java running in an application server or microservices architecture. The beauty of this approach is migration from one technology to the other. The code generation should mean that it is platform agnostic and that changes in technology do not require re-coding, merely re-deployment thus saving costs, time and risk.

This pragmatic three-layered approach that takes advantage of current technology could be used in the short to mid-term until fully machine-executable regulation comes along. Even in the long term the machine executable approach still leverages many of the same aspects of this approach, namely explicit terminology and orchestration that can feed into a model driven code approach.

Obviously, there will always be a certain amount of interpretation in terms of eligibility criteria because the legislation itself is always written with a bit of leeway. It's similar to Pareto's law - you're doing 80% automated and then 20% taking into account your own in-house legal counsel or business people and they might have a slightly different interpretation. If you implement the proposed three-tier structure with the important aspects of extensibility and custom logic at a modelling level, you benefit from consistency across asset class, across geography and more importantly, across all the different silos in your organization.

Auditability is another key advantage to this approach; if you're building technology such as Restful APIs, whether you're storing this reporting data in a blockchain, a distributed ledger or a traditional database, you would have the capability to store the audit trail and refer back to the report that links directly back to the legislation. Putting this audit trail in place would then help to avoid issues with fines in the future. Senior management would be able to prove why something was done and that it wasn't a deliberate case of incorrect or lack of reporting. Being able to evidence that something was caused by what was just a coding error could mean a reduced likelihood of a fine. If you can show the thought process you took in terms of reporting a transaction, the way it was reported, that you followed a logical and rigorous process to report it, a more lenient response from the regulator might be awarded. Clearly, there are benefits in being able to take the same approach across asset classes and across geographies and also show visibility of compliance to management.

Traditionally, compliance is not seen as a money-maker however, the approach to avoid fines could also lead to a minimal viable implementation approach. If, however, a business case is made clear, compliance can become an opportunity.

MAKING SENSE OF THE DATA LAKE

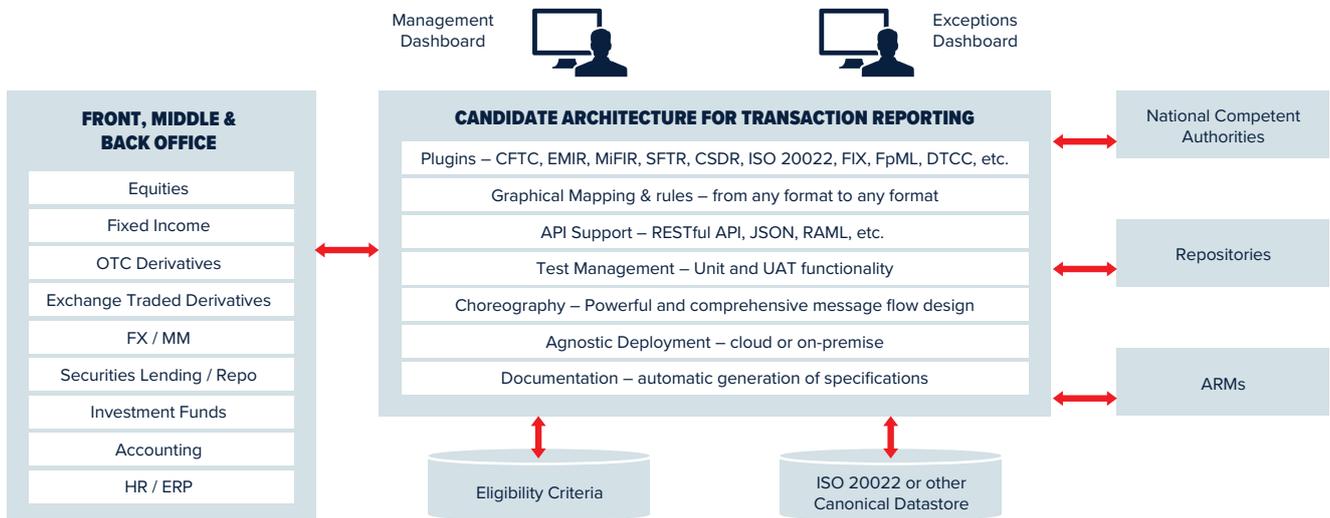
The ability to bring multiple different asset classes out of siloes into one place, into a store - a data lake - in which analytics can be performed to help you evaluate reporting requirements, over time can become a key tool that management can use to understand their organization and make better informed business decisions.

Many firms already have this information but it will often only reside in the front office. Transaction reporting by definition captures the whole lifecycle and thereby much more granular information about all aspects of trading through to settlement and reconciliation in the back office.

If you're funnelling everything through one big data store which contains evaluation criteria and eligibility criteria, then you're building up a huge record of what is being traded, who are the parties, etc. As this information has to be reported, why not put them all in one place? The technology is there to support that, whether it's in-house or in the cloud. Putting

those data stores and the execution environment together, whether it's in Amazon Web Services or Microsoft Azure or Google or whatever comes next, seems like a sensible architecture that would future proof as much as is possible.

The ideal solution would be something like the diagram below which brings together all of these key features and artefacts into a Regulatory Reporting Hub.



By adopting flexible, technology-resistant architecture and products in conjunction with externalized eligibility criteria and API capability to access other external reference data services, costs and timescales of compliance can be greatly reduced and visibility improved. Regulatory reporting will never be easy given the underlying complexity of the instruments being traded but by embracing the challenges and thinking outside asset class silos, a firm can at least reach a point that regardless of where regulation takes them in the future or wherever technology moves, you can adapt to that ever changing world without reinventing wheels.

ABOUT VOLANTE TECHNOLOGIES

Since 2001, Volante has been successfully addressing challenges around financial message integration and payments capture, pre-processing, processing and clearing.

We are committed to creating products that enhance the business agility that financial institutions depend on to remain competitive and thrive by quickly taking advantage of business opportunities as they emerge.

The inherent properties within all Volante products, including: extensive automation, a large maintained library of out-of-box plugins, transformations and processor modules, configuration rather than coding and inbuilt testing, promise significantly accelerated implementations and transformation projects at much lower costs than alternative solutions.

Supported by offices in Jersey City, London, Dubai, Mexico City, Bogota, Hyderabad, Chennai and Pune, Volante is able to encapsulate a best practice approach into all its product lines.

For further information please visit: www.volantetech.com

