



The Evolution of Buy vs. Build

From cloud to open-source, APIs, AI and beyond, tech evolution is changing how buy-side firms procure new tools.



CHARLES DARWIN

Customise your email alerts with Follow Topics

Our most popular Follow Topics include:

- Data management
- Artificial intelligence
- Management & Strategy
- Alternative data



Get started: www.waterstechnology.com/following

Editor-in-Chief Anthony Malakian
anthony.malakian@infopro-digital.com
Editor at Large Max Bowie
max.bowie@infopro-digital.com
European Editor Jo Wright
joanna.wright@infopro-digital.com
Editor, Asia Wei-Shen Wong
wei-shen.wong@infopro-digital.com
News Editor Josephine Gallagher
josephine.gallagher@infopro-digital.com
US Staff Writer Rebecca Natale
rebecca.natale@infopro-digital.com
Reporter Nyela Graham
nyela.graham@infopro-digital.com
Reporter Melissa Fleur Afshar
melissafleur.afshar@infopro-digital.com
Head of Editorial Operations Elina Patler
elina.patler@infopro-digital.com

Global Content Director Victor Anderson
victor.anderson@infopro-digital.com
+44 (0) 207 316 9090
Commercial Director Ince Saleem
Tel: +44 (0) 20 7316 9258
ince.saleem@infopro-digital.com
Business Development Executive Sonja Patillo
Tel: +1 212 776 8083
sonja.patillo@infopro-digital.com
Account Manager Daniel De-Bruce
Tel: +44 (0) 20 7316 9126
daniel.debruce@infopro-digital.com

Marketing Manager Louise Sheppey
tel: +44 (0) 20 7316 9476
louise.sheppey@infopro-digital.com
Design Lisa Ling

Corporate and Single Subscriptions
US: Barbara Falman tel +1 646 736 1852
info@waterstechnology.com

Global Brand Director Katie Palisoul
katie.palisoul@infopro-digital.com
Global Editorial Director Duncan Wood
duncan.wood@infopro-digital.com
Managing Director David Pagliaro
david.pagliaro@infopro-digital.com

Infopro Digital Head Office
Fifth Floor
133 Houndsditch
London EC3A 7BX
United Kingdom
tel: +44 (0) 20 7316 9000
fax: +44 (0) 20 7930 2238

Infopro Digital US
80 Pine Street, 21st Floor
New York, NY 10005
tel: +1 646 736 1888

Infopro Digital Asia
Unit 1704-05
Berkshire House, Taikoo Place
25 Westlands Road
Quarry Bay
Hong Kong
tel: +852 3411 4888

Infopro Digital Customer Services
tel: +44 (0) 1858 438 800

In order to receive *WatersTechnology* magazine every quarter you must have a *WatersTechnology* Subscription or a *Waters Premium* Subscription. For more information and subscription details, visit waterstechnology.com/subscribe

WatersTechnology (ISSN 1068-5863) is published quarterly (four times a year) by Infopro Digital Risk Limited. Printed in the UK by Stephens & George Print Group, Dowlais, Merthyr Tydfil, Wales.

Published by Infopro Digital Risk Limited. Copyright Infopro Digital Risk Limited (IP), 2022. All rights reserved. No part of this publication may be reproduced, stored in or introduced into any retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of the copyright owners.



The industry's increasingly cloudy skies

Recently, Rebecca Natale wrote about a new company started by Nick Kolba. He was a key architect on Thomson Reuters's Eikon platform, and he led the creation of the FDC3 protocol while at OpenFin. FDC3, as you may know, is the open standard used by those in the desktop app interoperability space.

For his next trick, Kolba is looking to move beyond the desktop and into the cloud. As Reb writes, his startup, Connectifi, will allow applications to "talk" to one another and work in sync, and users can build their own custom workflows. But in a significant departure from the way this is usually done, Connectifi will utilize the cloud (and allow for mobility) instead of a desktop container, such as Electron. From the story:

Kolba was inspired, in part, by the industry's push toward the cloud, which accelerated through the rise of remote working and by targeted offerings from Big Tech firms such as Amazon Web Services, Google Cloud, and Microsoft Azure.

"Interop benefits from web tech because it puts everything on an equal footing, and you have less fragmentation in the technology stack. But at the same time, the way it's implemented today means you end up with another fragmentation across different container platforms," Kolba says. "What I wanted to do was create interop that's not platform-dependent, not desktop-dependent, not container-dependent, not install-dependent; that can run anywhere; and that's truly native to the web and truly native to the cloud."

The ways in which financial technology is developed and delivered are rapidly changing (see page 52). The idea of desktop app interoperability is still in its relative infancy. But Kolba sees an opportunity to already move beyond this particular brand of tech development by leveraging the cloud for app interop. I fully expect the likes of OpenFin, Cosaic and Glue42 to move down this path in the near future (if they haven't begun already).

While the cloud is hardly new to us and our smartphone-dominated consumer world, the matchmaking between trading challenges and cloud solutions is. It's easy to write "The industry is rapidly moving to the cloud," and state it as a fact. I mean, it is a fact, but the statement is often lacking a tangible picture of that migration. But in true *Nineteen Eighty-Four* fashion, the problem is that if something is stated as fact often enough, whether true or not, it becomes fact.

An example: "Blockchain will revolutionize the capital markets!" Well, where's the proof? See what I'm getting at? When this happens—to steal from another famous book—senior technologists and data professionals start tilting at windmills.

Cloud is the revolution, and the proof can be found in the following pages, as most every story in some way touches on cloud delivery, storage or analytics. And I imagine that virtually all stories we write going forward will also at least touch on cloud in some way. It's not just the revolution—it's also the future. The question is: How will firms evolve to capitalize on the cloud? **wt**

Anthony Malakian
Editor-in-Chief

waterstechnology

Contents

1 **Editor's Letter**

4 **New Perspectives**

24 **Newsdesk**

25 **Open Outcry**

27 **Sell side ramps up outsourced trading desk services**

A raft of new entrants are offering outsourced trading services to buy-side firms, anticipating a wave of takeup among larger asset managers, driven by cost and coverage needs. But are they aiming too high?

By Max Bowie

32 **Beyond the 'big red blob'**

UBS has started a journey to implementing mesh data architecture, a new concept for federated data availability in enterprises aimed at getting value out of analytical data.

By Joanna Wight

36 **Retooling repo**

From brokers to trading platform providers, players in the repo market are focused on initiatives around reducing settlement fails, increasing automation, and streamlining operations. By Nyela Graham

40 **From pharma to finance: Cracking the DNA of data management**


Enterprise data management has traditionally addressed any aspect of financial data across an organization. But as investment firms' portfolios of enterprise data broadens in definition to include other types of non-financial data, EDM projects must also expand to embrace new data types. By Max Bowie

44 **Cloud and APIs begin to (slowly) permeate the post-trade space**

As financial firms turn their attention toward modernizing the back office, how they approach these projects comes under new scrutiny. By Wei-Shen Wong

52

COVER STORY:
How tech evolution has changed the world of finance
By Anthony Malakian




27

Sell side ramps up outsourced trading desk services
By Max Bowie

32

Beyond the 'big red blob'
By Joanna Wight



48 **After switch to Google, Symphony begins data migration for clients**

Symphony decided to first migrate client data to Google from Amazon Web Services, and then focus on applications. By Melissa Fleur Afshar

52 **COVER STORY: How tech evolution has changed the world of finance**

Neal Pawar, the former CTO of AQR and current COO of Qontigo, chats with Anthony Malakian about some of the major trends that are changing how asset managers interact with the vendor community, and how this shift mirrors the most significant evolutions in capital markets technology over the last decade.

58 **Human Capital**

60 **Opinion: The power of FDC3 appD: A universal standard to distribute and discover applications**

36

Retooling repo
By Nyela Graham




GET MORE ONLINE

News. Webcasts. Video. Audio. Special Reports. Get it all at waterstechnology.com

40

From pharma to finance: Cracking the DNA of data management
By Max Bowie



48

After switch to Google, Symphony begins data migration for clients
By Melissa Fleur Afshar



Finos will host common domain model for Isda, Isla and Icma

The non-profit will provide a neutral, open-source repository to help the trade associations govern and maintain the CDM. By [Josephine Gallagher](#) and [Rebecca Natale](#)

The Fintech Open Source Foundation (Finos) has been selected to provide an open-source repository for the common domain model (CDM), a single digital representation of trade products, events, and actions across a trade lifecycle.

“Following a request for proposal earlier this year, the International Swaps and Derivatives Association (Isda), the International Capital Market Association (Icma), and the International Securities Lending Association (Isla) are pleased to confirm that Finos has been selected to provide an open-source repository for the common domain model,” says a spokesperson for the three industry bodies. “We look forward to working with Finos on this important initiative, which will extend the benefits of the CDM from derivatives to securities, repo, and securities lending.”

The three industry associations signed an MoU in August 2021 to strengthen their collaboration on developing the CDM. As part of the joint initiative, they invited organizations to put forward service proposals that would meet their requirements for a neutral third-party host. The RFQ specified that it needed this host to provide a central repository in which the CDM user community could access and consume the model; to facilitate continuous development and maintenance of the CDM code; and to help the Isda, Isla, and Icma in governing the model and growing awareness of it in the financial services industry.

Finos, the umbrella body for financial services under the Linux Foundation, was a logical choice: The body promotes the adoption of open source and open standards, and has worked with Isda on



a pilot project using Legend, a modeling language that Goldman Sachs contributed to Finos, to develop extensions to the CDM.

The CDM, originally founded by Isda, began as a blueprint for how over-the-counter (OTC) derivatives can be traded and managed across the transaction lifecycle. It is a single digital representation of trade events and actions. The model has since expanded beyond derivatives to other product types. For instance, Icma, under its CDM Steering Committee, is applying the model to the repo and bond markets, and Isla’s CDM Working Group, in collaboration with tech partner REGnosys, has developed an early version for modeling securities lending transactions.

“If you think of those three product sets, there is a core CDM that underpins them. It could be a set of common data types, values, and definitions for what a quantity is, what a product is or what a trade is. These are the kinds of common things that we all use,” says Ian Sloyan, director of market infrastructure and technology at Isda.

Outsourcing the CDM governance will offload some of the logistical work currently managed by the trade associations. A simple example is that communications about the CDM for OTC derivatives are hosted and run on Isda’s internal systems. For instance, Isda’s mailbox, or Microsoft Teams accounts, are being used to coordinate working group calls for projects like its Digital Regulatory Reporting (DRR).

Sloyan tells *WatersTechnology* that another reason for using a third-party entity was to communicate to the industry that the CDM is not solely a derivatives solution but an open, free-to-use framework.

“We want to make a bigger statement about how the core CDM is open-source, and other trade associations, should they want to develop things for different markets, can similarly build on this neutral CDM,” he says.

Another core objective of the decision is to encourage adoption, participation, and contribution to the CDM code. “It’s hopefully opening up the tent, or at least building a bigger tent for us to progress,” Sloyan says.

Cracking the code

One core area that Isda’s CDM aims to resolve is trade reporting. In February 2021, the trade association launched its DRR initiative to help interpret and implement regulatory reporting rules consistently using a common, machine-executable code. The DRR has several working groups that interpret the rules, apply them to the CDM digital blueprint for the derivatives business, and translates those terms into coded instructions that are then used in reporting systems.

The objective of Isda's DRR is to help investment firms with the constant barrage of regulatory changes, putting pressure on compliance teams and forcing them to frequently update their internal systems. For example, the US Commodity Futures Trading Commission's (CFTC's) upcoming revision of swap data reporting rules will come into effect on December 5, making it the first regulator to amend its framework to incorporate globally harmonized, unique product identifiers and other critical data elements (CDEs) intended to improve the consistency and accuracy of the reported data.

Other jurisdictions, such as Europe and Asia, will follow in the US's footsteps to align their rulebooks by including new CDEs and end-to-end reporting in ISO 20022 XML, a messaging standard created by Swift and used for reporting to trade repositories or swap data repositories under the CFTC. The European Securities and Markets Authority (Esma) consulted on the draft changes to the European Markets Infrastructure Regulation, known as Emir Refit, in July 2021 and the rule updates will take effect 18 months after they are finalized.

While reporting firms will have to undergo the painful exercise of preparing for the new global rule changes, one silver lining is that there will be a significant overlap in the compliance work across jurisdictions. Together with that, Isda's DRR project aims to alleviate some of the compliance costs on firms.

"There's probably about 70% overlap between Emir and CFTC, so why try to solve the problem twice when you can solve it once? If there's a rule that has interpretation across Emir and CFTC, and both communities attend our peer review meetings, we can solve it just once," says Alan Milligan, Isda's head of data and digital solutions.

The Isda DRR initiative is made up of members from banks, asset managers, trade repositories, and vendors and is currently focused on codifying the CFTC rewrite. As part of the agile development project, there are two working groups.

One that meets monthly and is responsible for interpreting the CFTC rules. The second working group, which consists of 30 to 50 contributors, is focused on turning the interpretations into coded executable instructions. Milligan says that the groups have 98% of the coded instructions completed against the CFTC rules. On July 1, during a CFTC Peer Review, SDRs and reporting firms had the opportunity to test the DRR CFTC codifications but because this was done at their own discretion, it is unclear how many opted to use Isda's DRR.

In search of common ground

Technology is only one part of the compliance puzzle; understanding what the rules mean and how they should be understood is often a bigger challenge. If bank A interprets a rule one way and bank B interprets it a second way, this sets off a chain reaction of different tech design implementations, leading to different data fields being filled, and often inconsistent data being reported to their national competent authorities.

In other cases, banks might have similar interpretations of the written rules but still end up having inconsistent, and often costly, bespoke technical implementations, says Andrew Bayley, director of data reporting at Isda. Downstream this could also impact data quality and matching rates to their designated TR or SDR.

"They may well both be correct but because of the way the rule is written up, because of the way it's been interpreted, you may end up with two processes which aren't compatible," says Bayley.

The goal of DRR is to mutualize the work of interpreting the rules and create common implementations that lead to cost savings for the industry. This might sound great in theory, but it's a whole other beast putting it into action. One barrier to adoption of the CDM for reporting is creating confidence in the codified interpretations of the rules so that it meets the compliance needs of the regulators.

In parallel, some regulators like the Financial Conduct Authority

and the Bank of England are exploring their own pilot programs for developing machine-readable and executable regulation, which, ironically, is also named the Digital Regulatory Reporting (DRR).

To test Isda's DRR outputs, Bayley says the working groups cross-check the data against member-provided reported data and is engaging with regulators periodically to validate interpretations and outputs.

"The regulators are unlikely to say, 'Yes, this is 100% correct,' but they will point out to us if we've got it wrong," he adds.

Adopt or shop?

Developing and releasing a domain model in one thing, but its success, however, is contingent on its adoption. There are 30 to 50 contributors to Isda's DRR, but even the trade association is unclear as to how many of those are using the CDM for the codifications for their native reporting or as a validation tool to benchmark and reconcile their reporting framework.

On July 1, during a CFTC Peer Review, SDRs and reporting firms had the opportunity to test the DRR CFTC codifications but because this was done at their own discretion, it is also unclear as to how many opted to use Isda's DRR.

Bayley says adoption "is not all or nothing." The DRR codifications are open and free to use, and firms have the choice to use it as their primary or secondary reporting method, but Isda does not have precise numbers on firms using it internally.

In the past, however, the original CDM releases have shown patchy uptake, with some citing the model being too complex to use or some struggling to dedicate resources to the scale of work required to map their internal definitions to the CDM. Milligan says the cost of implementation is a personal one. In other words, each firm will have to decide if or how they want to adopt the CDM into their process and reporting framework. [WT](#)

Geolocation data market contractions highlight importance of cost, context

All is not well for providers of geolocation data, with some slashing staff or shutting down entirely. Those still thriving are the ones who realize it's no longer all about 'location, location, location.' By [Max Bowie](#)


Geolocation data—information on the movement of individuals or commodities sourced from mobile devices and transponders on shipping vessels—has emerged as a valuable dataset in investment firms' efforts to predict supply and demand and monitor consumer trends.

For example, if you can use individuals' cell phone signals—which are accurate to within a meter or two—to monitor foot traffic in retail stores and combine that with transaction data from credit card companies, then you can predict a company's sales and revenues in real time, rather than waiting for company updates, quarterly filings, or annual reports. If you see sales rising, you can buy company stock before it releases its official results, in anticipation of the stock price rising.

Yet the rising cost of compiling and maintaining databases of location-based information and combining it with other types of data is proving too much for some providers.

Earlier this year, Gyana, a UK-based startup location data provider turned no-code reporting platform founded by University of Oxford graduate students, closed down, while Denver, Colo.-based mobile location data provider SafeGraph was forced to lay off a quarter of its workforce.

Gyana and SafeGraph aren't the first location data providers to experience problems. In 2019, New York-based Thasos Group, one of the pioneers of selling location data to financial services firms, had to cut around two-thirds of its staff and its CEO stepped down, according to reports at the time. In March 2020, *WatersTechnology* reported



that the company's assets had been acquired by Long Island, NY-based Market Service, Inc., which owned various data and business intelligence brands, including AggData—also a provider of location data—and CreditIntell. F&D Reports and Arms, which specialize in credit risk reports for retailers and privately held companies.

“For any type of data business, it's very difficult if you are selling into one market. But we're a very diversified business selling into many sectors outside financial services,” says AggData managing director Josh Suffin. Since the acquisition, Market Service has integrated the Thasos assets across its brands, built a new dashboard for accessing the data, enhanced Thasos' original algorithms, and integrated more location data from AggData.

Ironically, Gyana's no-code platform was designed to help users across industries from private equity to real estate and technology companies to gain insights from mobile data without needing their own data science organization.

“In many ways, Gyana has never been in such a good position. We have a product our users love, and an effective, remote-first culture with minimal bureaucracy and lots of getting shit done. Unfortunately, the numbers just don't add up. Gyana is an expensive product to scale. We don't have the capital, and thanks to founder dilution and difficult market conditions, a fundraiser isn't feasible,” said co-founder and CEO David Kell, explaining the decision to shut down in a blog post.

SafeGraph's CEO Auren Hoffman described his company's decision to cut about 25% of company headcount as a “difficult yet important decision to lower the cash burn,” and conserve cash. “We made this decision even though we had two years of cash in the bank before the cuts. Because in this environment, two years is not enough,” he said in a post on SafeGraph's website.

But while Suffin says AggData has been able to benefit from location data, he notes that not all providers are experiencing success. “Currently, in the market, we are seeing contraction and companies going away because it's a costly business to be in, and sales are difficult and—particularly in financial services—can take a long time because of these firms' compliance requirements,” he says. The key, he says, is to be well-funded, not burn through cash, and not over-hire before having the revenues to support expansion.

“Thasos was doing a very good job with data quality and coverage, they had a great product, and if they'd had more funding, they'd probably still be around today,” he adds.

Mounting costs

Yiannis Tsiounis, CEO of New York-based geolocation data provider Advan Research, says location data has become a much more expensive proposition since Advan was founded in 2015—and providers must be able to offer better service to stand out in a crowded market.

“Because there’s more competition, you need more points of interest, and your data must be cleaner. If you want to compete, your data must be of better quality and have greater coverage. So, it’s harder to start from scratch, because if you want to compete, you need to have a solution immediately.”

Thanh-Long Huynh, CEO of Paris-based macroeconomic analysis provider QuantCube, agrees. Beyond simply assembling a comprehensive dataset and augmenting it with related data, there’s a lot of processing that vendors must perform on the data before it’s fit for use by clients.

“It’s taken almost two years for our team to clean everything in a systematic way,” Huynh says of QuantCube’s shipping-related location data and all the information that accompanies it, adding that clients wanting to use the vendor’s data simply don’t have the resources or time to process the data to make it fit for use.

Likewise, he says location data derived from mobile devices is “very powerful” but also very difficult to access and process. “It took us two-and-a-half years before we got the first results. But now that we’re over that entry barrier, I think there are many applications that could be derived from that data,” Huynh adds.

The effort needed to match that level of data within a short timeframe would normally require a large capital investment of anywhere between \$20 million and \$200 million, Tsiounis estimates, whereas Advan was able to build its full offering over time.

But keeping on target is hard when others are constantly moving the goalposts. For example, as demand increases, the underlying data sources realize their

data has value, so they charge more for it, Tsiounis says.

In addition, location data on its own has little value: It’s when that data is combined with other information that it becomes valuable. Location data in stores doesn’t tell you whether customers are buying goods, or how much they’re spending: Point-of-sale transactions data does that. Likewise, the location of a ship on its own says nothing. But knowing what kind of vessel it is, what it carries, and using shipping manifest data to understand what it’s carrying, and where it’s sailing to and from, creates a much richer picture—and one that can be used to make assumptions about oil prices, based on supply and demand.

In QuantCube’s case, the vendor collects more than 14 billion data points that encompass much more than just location data, many of which augment that data and enable users to draw more accurate conclusions—for example, in addition to tracking 80,000 vessels and updating their position every 15 seconds, the vendor also captures atmospheric data from satellite photographs, and monitors social media in different languages most relevant to specific commodities. For instance, social media posts in Arabic contain five times the number of oil-related posts as English-language social media, Huynh says.

Those extra datasets come with a price tag—not to mention the cost of a team of data scientists to marry the datasets and derive insight that delivers value for clients. And that’s also a moving target: As vendors and their clients identify new use cases for the data, vendors need to license new datasets to support those uses. For example, using mobile device data to monitor foot traffic at a residential building can be a real-time indicator of occupancy rates, which can serve as an input to pricing commercial mortgage-backed securities more accurately.

Another use could be using data related to the mobile devices—within the confines of the General Data Protection Regulation (GDPR)—such as whether that device’s owner (an individual or a

corporation) has paid their monthly cell phone bill or not, to use as an input to credit risk assessments, Huynh says. Or, by analyzing the mobility of company-issued mobile phones, an investor could expect to see higher or lower revenues at that company based on how mobile its staff are.

And as more datasets emerge, these in turn create the potential for new use cases. “There are certainly lots of alternative data points being created by the internet of things/wearables, app installs and geolocation, but they are not as widely used as other ‘alt data’ sets,” says David Easthope, head of fintech in research firm Coalition Greenwich’s market structure and technology team. And while noting that “geolocation or footfall is more of a specialized dataset used fairly selectively by certain hedge funds,” he also adds that geolocation data often appears on the future roadmaps of funds not currently using it.

In the timeframe of those future roadmaps—possibly within three to five years—QuantCube will be able to estimate revenues and financial statements for companies in real time using combinations of alternative data, Huynh says.

Context is king

However, location data needs to be consumed in conjunction with those other datasets to have value. “Location data is only one part of a whole lot of analytics we need to analyze macroeconomic trends or international trade. It’s extremely powerful, but not just by itself,” Huynh says, warning that the market will not be large enough to support a multitude of pure location data providers, and that vendors will need to diversify to provide a full data offering, of which location data is one component.

And key to getting the most out of that data component is being able to understand properly how location data can be used. Once a firm or vendor fully understands how to view information provided about and by a location, they can grasp how to develop any number of potential analyses and use cases. 

The newest trade repository in town preps for uphill battle

The ex-head of CME Group's repository services looks to disrupt regtech with his startup, recruiting talent from market structure firms that have scaled back regulatory reporting businesses. By [Josephine Gallagher](#)

Where some have failed, others hope to succeed. Jonathan Thursby, CEO and founder of startup trade repository and regtech provider Kor Financial, is hoping that he can succeed in an industry where large incumbents have been winding down their businesses.

As regimes like the European Market Infrastructure Regulation (Emir) brought in post-financial crisis transparency rules, large market infrastructure providers leveraged their capabilities to offer trade repository and trade and transaction reporting services. But over the past two or three years, these businesses started closing. In 2020, CME Group scaled back its regulatory reporting and TR services. Deutsche Börse exited the market the following summer. UnaVista, the matching, reconciliation, and regulatory reporting platform of the London Stock Exchange, shuttered its Securities Financing Transaction Regulation unit in January of this year.

These businesses found that trade reporting is a tough area in which to thrive, due to costly overheads and having to keep up with the relentless cadence of regulatory change. For businesses that offer trade repositories and swap data repositories (SDRs) in the US, there is an additional layer of cost: as highly regulated entities, they face penalties for reporting infringements. The European Securities and Markets Authority (Esm), for example, has slapped the DDRL trade repository of the Depository Trust and Clearing Corp. (DTCC), UnaVista and EU trade repository RegisTR with fines over the past two years.

This is apparently a harsh landscape in which to be launching a new



business, but Thursby, who headed up CME's repository services and EU TR before the businesses were wound down in November 2020, believes that the incumbents exiting the sector have created a vacuum for rival companies to fill. "This is a very skilled and specialized area, and you must deeply understand what the problems are and come up with ideas of how to solve those—that was the genesis of Kor," he says.

Thursby launched Kor about two years ago to provide a reporting-as-a-service offering for over-the-counter (OTC) derivatives, and as a swap data repository under the Commodity Futures Trading Commission (CFTC).

He has since stacked Kor's C-suite with expertise from the reg reporting firms that scaled back their businesses. Tara Manuel, the former head of Ice's Trade Vault, is Kor's chief compliance officer; Tom Wiczorek, former managing director of products for regulatory reporting and post-trade operations platform at UnaVista, is chief product officer; and Rahul Cherukumalli, CME's former head of regulatory reporting, has been brought in as director of client success.

In the future, the vendor has ambitions to create a marketplace for regulatory reporting services that would work in a way analogous to Apple's iOS app store.

But for now, Kor's focus is on its SDR, which was registered with the CFTC in March—the first such approval granted by the commission since 2014. The startup announced in June that it had completed a Series A funding round led by venture capital firm Mosaik Partners, to accelerate the deployment of the SDR and kick-start expansion efforts. Kor planned to go live with the SDR by early September with support for the CFTC rewrite changes. The number of clients it will launch with is in the teens.

Thursby says clients have been able to test the changes to the CFTC reporting framework for several months, to help prevent any switchovers forming bottlenecks ahead of the first phase of the CFTC go-live at the end of the year.

As the company grows, Kor will look to become a trade repository in other jurisdictions and build out a technical suite for servicing a long list of regulatory reporting needs. Thursby anticipates that another application to become a security-based swap data trade repository with the Securities and Exchange Commission (SEC) will be approved soon. He says he expects Kor to be registered as a TR with both European and Asia-Pacific regulators.

Once its registration process is out of the way, Kor will later move on to the second phase of its expansion. The company intends to develop and offer regulatory reporting services that will complement the TR endpoints but will operate as separate businesses.

“In terms of the trade repository itself, you want that to be fairly straightforward and well defined,” Thursby says. “It’s tough to offer additional value-added services under the regulatory purview, and so you operate those separately.”

While Kor’s executives have big plans, the first step is convincing trading firms that a startup can do the job of reg reporting better than the stalwart incumbents—essentially, they’re looking to prove that disruption is possible in this challenging marketplace.

Thursby recalls during his time at CME running a unit starved of resources and investment—a common complaint of many working in post-trade functions in all kinds of companies across capital markets. “It became this dreaded experience and you always felt like you’re at the other end of the club. None of that has changed in the last 10 years—there really hasn’t been any evolution [for trade reporting],” Thursby says.

But firms might have to invest in post-trade now, as a second wave of regulatory action on transparency regimes drives significant demand for reporting services. With regimes like the EU’s Emir and the CFTC’s Dodd-Frank Act now firmly established, regulators are getting to the point of reviewing them to improve the quality of reported data and harmonize regulatory rule books. Emir is going through a Refit; the CFTC is rewriting the reporting rules for OTC derivatives.

“Initially, the expectation [of reporting entities] from regulators was getting the information to them,” says Vinod Jain, a strategic advisor at consultancy Aite-Novarica. “The focus is now to improve the data quality, consistency, completeness, and the clarity of the data.”

Building blocks

In the reg reporting space, scale is important to offset overheads and run a sustainable business. Kor won’t have size on its side, at least not to begin with. But Thursby believes that by starting a company from scratch, and by using tools like the cloud that were either non-existent

or too nascent for live production a decade ago, his startup will have the advantage of agility, onboarding new customers quickly.

He says Kor’s cloud-native platform was built with the intent of resolving data problems first, and trade reporting second. For example, reporting firms are having to reconfigure their systems and translate their data to fit schemas like ISO 20022, a standard for financial data exchange.

“The platform we’ve built is more powerful, in our opinion, than what trade reporting actually needs, because we were trying to solve a more fundamental issue in the industry about data management and the evolution of schemas over time.”

Kor’s technology shares some qualities with distributed ledger technology, but one key difference is that Kor is a centralized platform. It is an event-driven system: It always houses one version of the data, and the data is immutable.

“Everything that comes onto our platform is an event. One event can trigger another event, can trigger a third event, can trigger a fourth event, and we store all those events. We always have an instant audit trail of everything that ever occurred and all the metadata related to it, and we store all of that infinitely,” Thursby says.

Thursby, who spent 13 years at CME, says most large enterprises will spend 30% to 35% of their technology resources on maintenance, re-engineering, and patching systems. But there are ways to design systems to significantly cut the cost of operations and upkeep.

“Most of our competitors today run very separate applications, very separate installations, with very separate code bases. We get a lot of efficiencies from effectively having a common code base across everything that we build. And the way that we design the software into small services means that it’s much more supportable,” he says.

David and Goliath

The trade repository market is still dominated by incumbents like the

DTCC, RegisTR, and UnaVista. And many companies, both large and small, specialist and generalist, exist in the regulatory reporting space. But observers say there is still a growing need for these services.

Virginie O’Shea, founder and CEO of consultancy Firebrand Research, says there is a need for more services in the space that Kor is trying to disrupt. Many vendors in the regulatory reporting market deal with initial regulatory reporting tasks, such as cleaning and managing reported data.

However, a startup like Kor would have to be able to convince banks and asset managers that its business is sustainable, as firms have little appetite for building new connections to yet another TR. This work is a painstaking process that involves migrating the firm’s historic data to the new provider.

“That would be the investment, and even if it’s just time—sometimes it isn’t a massive amount of cost—but it could be the time required by operations or the IT team to look at these things and do their due diligence. So, there must be a compelling argument to do that,” she says.

As several large trade repositories and reg reporting providers have bowed out over the last three years, end-users will need assurances that Kor will still be around in a year’s time, O’Shea says.

Aite-Novarica’s Jain says there is pressure on reporting firms to resolve exceptions and data breaks sooner in the reporting lifecycle. To do this, investment firms need better analytics, dashboards, and visibility from their service providers to be able to compare their side of the trade with that of their counterparties.

Jain says Kor has the core infrastructure necessary to run such a business, and it has the in-house experience to run a trade repository. But for any new reporting firm to succeed in this next phase of regulatory reporting, they will need to have a sophisticated analytics proposition—one that keeps both the client and the regulator happy.

“It will all depend on the user experience,” he says. [WT](#)

In finance, all APIs are not created equal

With the rise of APIs, it's easier than ever for companies to position themselves as open to collaboration and eager to provide access, but some of these access points tell a different story under the hood.

By [Rebecca Natale](#)

By itself, the humble API is many things, none of which include “new,” “innovative,” or, often, even “interesting.” But it has underpinned a boom in technologies that *can* be described as such, thanks to its versatility: It can be used as a data delivery mechanism that gives firms greater control over how they access ever-increasing volumes of datasets, or bridge the chasms between disparate applications, at a time where connectivity, interoperability, and consolidation are market demands.


APIs have been around for decades, and they're relatively cheap and quick to create, making them an accessible tool for small startups and small trading shops alike—at least, that's the perception many in the tech community have.

But there's a difference between a working API by technical standards and an API that works for its users, a distinction not totally made by all those who build and use them.

“If you build an API, it's a highly complicated thing to do, and everybody thinks it's fairly straightforward,” says Thomas McHugh, CEO of Finbourne Technology, a buy-side technology provider established in 2016. “And what you end up with is people getting confused about what an API is supposed to be versus what a function is supposed to do.”

Finbourne offers what is essentially a data fabric layer that sits in the middle of other data providers and internal systems and translates that data across different languages, protocols, and terminologies. Underneath sits a data aggregation engine, and in the middle, an entitlements engine. On top is a data visualization layer.

The company builds the interfaces to other providers like Salesforce,



Snowflake, or Google BigQuery, so that when a user enters a command into Finbourne's query engine, it knows how to communicate with the other platforms' systems.

McHugh says that even though APIs are ubiquitous, he would like to see the industry achieve a greater understanding of business intelligence when building the next generation of them—for example, by creating associated ontologies so that different machines working with the same data can know what one another need.

(Semi)-intelligent design

While interoperability is a generic term—computer systems or software exchanging information—it's perhaps most interesting currently in the context of a subset of financial technology known as desktop application interoperability (though this name is misleading these days, as those working on desktop interop are expanding

to incorporate back-end and mobile capabilities).

In desktop interop, container technologies such as Electron wrap desktop applications into a package, in which each app can share context and information via APIs tailored to FDC3 standard, an open set of codified API specifications under the governance of the Fintech Open Source Foundation (Finos). Ongoing development of the standard, which just recently unveiled its version 2.0, is contributed to by desktop interop platform providers such as Glue42, Cosaic, and OpenFin, as well as any interested parties inside and out of financial technology.

Interop and FDC3 have been well received since their inception—notably by the sell side, among which Citi, State Street, UBS, NatWest, and Bank of America are adopters—but it hasn't been the sea change that some in the industry have hoped, says Glue42's head of product, Reena Raichura.

Raichura says the current desktop interop offered by providers is still only at a basic level of connectivity, partially because vendors limit what vendors like Glue42, Cosaic, and OpenFin can do within the offerings. Traditional companies that already reach thousands of customers aren't necessarily incentivized to open their products, and because the applications that they *can* connect for users are developed as "islands" in isolation, the workflows between them don't necessarily become intuitive or efficient.

"One of my frustrations now is with vendors themselves. Because they're halfway houses—either they're closed or they're partially open—we're not really getting the traction that we should be getting for desktop interop," she says. "Vendors say they have APIs, but when you actually speak to a user, and the user wants to go from A to B, or they want to do something specific, it might not actually be available in the API."

Raichura says many historically closed-off vendors have opened up through APIs to meet demand for collaboration and integration, but too often don't build them with a user's workflows in mind.

Her main objective is for those who build and use APIs to achieve "straight-through workflows"—a play on straight-through processing—which, in a literal sense, means cross-application workflow automation on the desktop. But for Raichura, it's much more about an overhaul of the design process in the way enterprise technology is built.

"Let's blur the boundaries of applications and look at what the business is doing, what the user is doing, and analyze those end-to-end business processes and workflows first before building a solution. That's when you're going to get a solution that actually matches the needs of a business and has real business impact. I think that's where the mismatch is today," she says.

Never change

Desktop applications and interoperability have become nearly synonymous in recent years, but interoperability—and therefore APIs—extends well beyond apps. It also goes beyond equities trading, the most technologically sophisticated asset class, to those that are still playing catch-up, such as fixed income.

Founded in 2018, MultiLynq is a startup that uses one API to connect bond market participants to the growing realm of electronic trading venues. Its service connects to 10 electronic trading venues—including Tradeweb, MarketAxess, Intercontinental Exchange's BondPoint, MTS BondPro, and others—via their own APIs, and allows clients to connect to all marketplaces via MultiLynq's single point of access.

Co-founder Patrick Scheideler says that the time it takes for a client to connect to any single venue's API on its own takes anywhere from three to six months, and the process includes a variety of hurdles. Different venues use different Fix messaging to communicate similar material such as "Where do you send your quote?" vs. "Where do you get the order?"

These small but specific specifications task firms' trading infrastructure managers with trying to identify how to make their own internal platforms—and their own configurations—work inside a new venue. Often, that scenario plays out many times across many venues for a single firm.

"The simple thing [would be] changing the messages, right? But that only gets you 5% of the way there," Scheideler says. "There's a lot more that goes into it. What does the workflow look like? Where are the edge cases? There's just a lot of time that goes into it."

The specifications end up being very long, very confusing, and still aren't always comprehensive—the only way to find out what's under the

hood, he says, is to do the integration. And that's not the fault of venues nor firms. Any spec is difficult to maintain, and multi-party connectivity, by definition, requires a lot of inputs, resulting in sometimes huge numbers of permutations that can't be covered in a single document.

It's the same conundrum that saw the advent of Fix and FDC3, both of which make communication between systems easier, but not truly easy—a side effect Scheideler attributes to an inability, or at least unwillingness, to go backward. Because APIs are often years and decades old, countless technology layers and specifications have been built over their once simple foundations.

Even poorly built or particularly old APIs rarely see updating, much less overhauling, says Finbourne's McHugh. The ones he'd classify as either are those not built according to frameworks that would make them efficient, such as the OpenAPI Specification, a standard interface to Rest APIs which allows humans and computers to discover and understand service capabilities without access to source code, documentation, or through network traffic inspection. The standard began as an open source initiative in 2015 under the Linux Foundation, led by giants such as Google, IBM, Microsoft, and others.

For example, McHugh says, if a vendor ships a software kit in C# or C++, but the end-user wants to use Go, that user would have to spend weeks or months building their own client for that unless the provider's product is compliant with OpenAPI or a similar framework. But if done the right way, APIs can be left on their own, he says.

"If they write an API, they are *never* changing it once it hits production status because they get that foundation correct," he says. "They have 10-, 15-year-old APIs, and they will not change it. And that's part of the contract. If you have a look publicly at our API, we have six, seven, 800 of them and if it hits production and enough customers have used it, we are never changing it." **wt**

Project Octopus becomes Octaura

Backed by the banks in the Project Octopus consortium, the new, independent company will launch this year with a focus on new trading protocols and integrated data analytics. At the same time, Bank of America will sunset its single-dealer loan trading platform, and Citi Velocity will end use of its BWIC protocol. By [Nyela Graham](#)

Last spring, *WatersTechnology* first reported that a consortium of banks led by Citi and Bank of America was preparing to combine its members' collateralized loan obligation (CLO) trading efforts into a new multi-bank trading platform under the working title Project Octopus. In June, the efforts of that consortium came to market as Octaura.

The independent company plans to roll out an electronic trading platform this year, with functionality for syndicated loans coming first.

Citi and Bank of America were joined by Credit Suisse, Goldman Sachs, JP Morgan, Morgan Stanley, Wells Fargo, and Moody's Analytics in the development of the platform, as well as the project's tech partner, Genesis Global, a low-code provider for financial markets. The group's efforts follow the trend of institutions looking to digitize the fixed-income market, especially loan trading.

As a result of Octaura's forthcoming platform launch, Bank of America will sunset its single-dealer platform, Instinct Loan Match. Citi will end the use of the bids wanted in competition (BWIC) protocol on its internal client platform, CitiVelocity.

Octaura will roll out syndicated loan functionality first, followed by CLOs. A gradual rollout is an effort not to overwhelm the market, says Brian Bejile, Octaura's CEO. Most loan trading is still done over the phone, and participants on the buy and sell side will need time to adjust their workflows, he says.

Octaura will also roll out an app marketplace for datasets and analytics to solve workflow needs. Moody's



Analytics, a partner in the Citi Velocity ecosystem, had already integrated its data and analytics into Citi's internal client platform, providing Bejile with a framework for an integration with Octaura. Data and analytics tools developed by other third parties will eventually be made available to clients and dealers through the marketplace.

The convenience Octaura aims to bring clients is modeled after ridesharing apps like Uber and Lyft. "They didn't invent transportation, but they made it convenient and available for us," Bejile says. "That's how I think about interoperability. What did Uber do? They took Google Maps, integrated with a payment system, and linked with companies that have existed for a long time."

For interoperability, Genesis Global's low-code platform allows for APIs that can handle inbound events or transactions and outbound data. Octaura is in the process of integrating with other

capability providers and hopes to present the final product in a unified, cohesive format, Bejile says.

Genesis's initial build was shown to more than 400 buy-side participants. In it, two protocols were used: Citi's CLO BWIC and Bank of America's Loan Match protocol. BofA's protocol allowed the bank to set the mid-market prices for a list of loans. During a match session, buy-side participants expressed buy or sell size interest through the platform. When two clients presented offsetting buy and sell interest or a match on a particular loan, a trade was executed automatically between those two counterparties.

The protocols act as daily scheduled events. "The question the client then asks is, 'What do I do between some of these auctions?'" Bejile says. The feedback indicated a need to create protocols that addressed how people were trading beyond these protocols. This created the inventory protocol.

"Dealer inventory is something that is unscheduled, it happens any time. Banks are advertising, 'Hey, I can buy this, I can sell that,'" Bejile says. "Clients can come in anytime during the day and engage the banks in the inventory."

The protocol is like a request-for-quote (RFQ), but there are two key differences. In an inventory protocol, the dealer starts the process by their listing sell and buy interest, while in an RFQ, a buy-side participant initiates the process by stating either buy or sell interest and size. The dealer also states the prices at which they are willing to buy or sell. A client in RFQ does not list a price, but rather requests a price from the dealers.

When Octaura goes live later this year, the Citi CLO BWIC protocol on

Citi Velocity and BofA's loan platform will be turned off. "Octaura's protocols make obsolete those predecessors by definition, so they will be sunsetted as part of Octaura going live," Bejile says.

In the making

Octaura began as an internal project at Citi, which Bejile headed up. He had previously spent 18 years at the bank and was global head of CLO issuer management when the seeds of Octaura were sown. Bejile commissioned an auction system for Citi's internal client platform, Citi Velocity, that would make the bidding process in CLO auctions more efficient than the usual method of picking up the phone and calling into auctions. The internal project undertaken by Citi allowed clients to submit bids electronically.

"We saw a 50% spike in the number of bids submitted to Citi by the buy side during the first week of launch of the feature," Bejile says. "It was clear that the functionality was both intuitive and much needed in the space due to that significant response."

Citi's senior management told Bejile to find another bank willing to partner on a multi-dealer platform. He reached out to Bank of America, which operates a single-dealer platform called Instinct Loan Match. The two banks joined forces, with Citi bringing its CLO expertise and Bank of America offering its experience in loan trading and platform building.

A June 2021 report from Coalition Greenwich found that e-trading in the syndicated loan market was growing after 18 months of an "e-trading tailwind" in the fixed-income market, spurred by office and trading floor closures caused by Covid 19.

Audrey Blater, a senior analyst in the market structure and technology group at Coalition Greenwich, says consortium support is uniquely important to the success of new fixed-income trading platforms unlike in equities markets.

"You need the backing of the dealer community," Blater says. "If they have skin in the game, so to speak, they're

going to be more willing to make prices, make markets, and use their networks to find the other side of the trade."

Non-consortium platforms may have great technology but never get enough clients and dealers to make an impact in the market, she adds.

Sources told *WatersTechnology* last April that in the case of Octopus, the move was a bid to prevent existing fixed-income platforms, Tradeweb and MarketAxess, from cornering the nascent electronic CLO market and charging what the banks see as high fees for trading and market data.

Bejile says the strategy was to start small, as opposed to consortiums that have started with multiple founding members. Today's leading fixed-income trading platforms were also born from consortiums. Tradeweb was founded in 1997 with investment from four banks. MarketAxess was formed in 2000 with support from JP Morgan and other large financial institutions, including Bear Stearns. Eventually, the banks ceded their investments. Thomson Corp. bought Tradeweb in 2004 and sold stakes to 10 major dealers in 2008. Tradeweb went public in 2019. MarketAxess went public in 2004.

Need for speed

A missing piece in the project's physical iteration was a technology partner. Citi enlisted Genesis Global, in which the bank had made a strategic investment in 2020 for an undisclosed amount. Citi's directive to Genesis was to build a multi-dealer platform that could support broad industry participation.

Genesis CEO Stephen Murphy says the provider was able to build a proof-of-concept in two weeks and a minimum viable product in three months. Bejile says the swift nature of building using low-code modules and tools made the vendor stand out among other tech providers the bank considered.

Proponents of low- and no-code software say that while such solutions certainly contain code, the user need not bother with it, as apps are primarily

built visually, or "Lego-style," by custom dragging and dropping pre-configured components and widgets onto a blank slate.

The appeal of low-code and no-code development is its advantages in building speed. Every new programming language is built, to some degree, on top of older ones. In programming's beginnings, there was machine code (binary ones and zeros) and assembly code (a low-level, plain-text version of machine code). These are most easily digested by a processor. An app written in a high-level code—say, Java—will get compiled down to machine code to run on a processor. No-code and low-code are extensions of this, which creates an additional level of abstraction that makes software building easier, faster, and more accessible to those with non-programming backgrounds.

Octaura isn't Genesis's only project in fixed income. The provider recently announced a new, expanded partnership with Neptune Networks, a fixed-income pre-trade market utility, on data distribution.


Murphy says Genesis has built generic technical components that can be applied across asset classes. In equities and foreign exchange, data needs to be distributed with low latency, so the provider built a component that could be configured to serve either of them. That same component was applied to the loan platform. Other generic components include authentication authorization and a transaction handler.

"If you log onto a system, it will look at what permissions you are allowed to see, if you've been authenticated, how are you authorized to do anything when you're on the system," Murphy says of the authentication authorization module. He says there are more than 50 of these kinds of components.

Octaura, through its name and logo, pays homage to its earlier name, Project Octopus. The company's new logo features an infinity symbol, representing interoperable workflows and data. But when turned on its side, the logo resembles an octopus. [WT](#)

FDC3 updates interop standard with 2.0 rollout as adoption grows

Finos gets ahead of FDC3's critical mass by making breaking changes to the desktop interop standard, including baked-in data exchange. By [Jo Wright](#)



In mid-July, at the Open Source in Finance Forum (OSFF) in London, Finos announced a new version of the FDC3 protocol, ticking over its mileage counter from 1.2 to 2.0.

Kris West, principal engineer at Cosaic and lead maintainer of FDC3, told the audience that as a single-API standard, FDC3 has been working as intended to get apps to talk to each other, speeding up the deployment cycle. Users are seeing the benefits of desktop agents and applications working in standardized ways, allowing users to rebuild their desktops, and drag and drop apps as needed, thus avoiding vendor lock-in and the single-use customizations that slow the cycle of acquiring and deploying an application.

“You might ask, then, if FDC3 is so successful, why did it need major

revision?” West told the audience. “We are rolling that first digit of the version number, which is conventionally done when you are making breaking changes to something. It’s something that, as a software vendor, you try to do very, very carefully.”

FDC3 is an open standard under the guidance of the Fintech Open Source Foundation (Finos), the financial services tech umbrella of the Linux Foundation, and development of the standard continues amid increasing adoption at financial services firms.

Of course, Finos wants FDC3 to be ubiquitous on financial services desktops, but ubiquity also brings challenges. More users from the vendor community, the sell side, and the buy side means widening user requirements.

“Hence we needed to seize the opportunity to change things before it got too difficult,” West said.

The four pillars of the FDC3 standard are API specs, a Rest-based application directory (or AppD), and constructs called intents and contexts. These four pillars are complementary. For example, a trader at a buy-side firm could use the vendor’s desktop agent to contact the AppD to discover available applications. Because of the standardized APIs, these various vendor apps can be slotted into or out of the desktop with ease, so users can always go with the app that meets their needs best.

FDC3 intents are often described as “verbs”—they specify what actions apps can execute, such as starting a chat or searching a blotter for an instru-

ment—and the contexts are like nouns, specifying within what data structure the intents can work.

The new version made many major refinements to the API specs, AppD, and the intents and contexts. The most impactful, however, according to West, is that the system of raising intents is now more of a two-way street. Earlier versions of FDC3 made it easy for a user to send requests to apps by raising an intent in their desktop agent, but they didn't get any information back. For example, a user could create an order, but would not get back details about the order once it was filled.

That kind of data exchange is now baked into FDC3, West said.

“Essentially, you can now automate any form of ‘Crud’ (create, read, update, and delete) operations you need to do from one application to another. You can send it over there, get back the update—you can have the updated record back. Then you can carry on your workflow with that data,” he said. “You can also create APIs to deliver datafeeds, such as pricing streams, streams of trades, streams of orders, details, updates, and customer activity that can then be handed on to other applications.”

Growth spurt

In 2018, OpenFin and industry participants unveiled the FDC3 protocol—a set of codified specifications that aims to enable traders' desktop applications to share information and interoperate—under an Apache 2.0 open-source license.

Vendors, notably OpenFin, Glue42, and Cosaic, emerged to provide interop platforms, known as desktop agents, and some firms are building their own desktop agents internally, with FDC3 as the standard. Consultancies like Adaptive Financial Consulting and Scott Logic are building tech on top of these platforms and using FDC3 in implementations with buy- and sell-side firms, seeing oppor-

tunity in the market that involvement in the development of FDC3 affords them.

And adoption among financial firms is increasing, primarily on the sell side, with firms like Citi, UBS, State Street, NatWest, and Bank of America all officially using tech built on FDC3.

Another sign that FDC3 has evolved is that new projects are growing up in its ecosystem. Nick Kolba, a veteran capital markets technologist and an originator of FDC3, contributed a project to Finos earlier this year. The project, called Electron FDC3, offers developers a neutral reference point and the full stack of FDC3 implementation, including a desktop agent and an application directory. (Kolba has since formed his own company that looks to take interoperability out of the desktop and put it into the cloud.)

As FDC3 sees more uptake, Finos has received a lot more user feedback to work with. This feedback has formed the backbone for improvement of the standard, which mostly takes place through frequent and regular discussion groups, and led to the recent update of FDC3.

Riko Eksteen, head of desktop strategy at Adaptive and co-lead maintainer, said that users were asking questions like, “How can I have intents that return data back?”, “How can I target a particular app instance?”, “What context and intents should I use—there aren't enough?”, “How can I publish my app easily?” or, “How can I make it work across different containers?”

Eksteen, West and others shared the view that the standard was confusing in having four specs that were also complementary, and some of the APIs needed clearing up, Eksteen added.

“How can we ensure each desktop agent will work the same? We need more examples and tools!” Eksteen said. “These were all things we were hearing from people, and it was daunting to figure out how to cover all of them. But that's where the discussion groups came in.”

Hammer, meet nail

Future improvements to the standard will now be additive—rather than breaking changes—with, for example, more intents available. For some in the FDC3 community, this is a good time to broaden the conversation away from just the technical specs of the standard.

Adaptive CEO Matt Barrett tells *WatersTechnology* that when the project began, it was in the hands of technologists who were very much getting into the weeds of the specs. They were followed by the vendors that built out capabilities and products based on FDC3. But the next stage of the discussion should be around the business benefits of FDC3 and the outcomes for users, rather than the technology itself, he says.

“Technologists tend to talk about the tool: ‘Look at my shiny hammer—isn't that a great hammer? What would you like me to do with this hammer?’ But I don't want a hammer; I want to put a nail into a wall. This is not a criticism; this is a natural phase of projects like this, and FDC3 is well-placed to do the next phase,” Barrett says.

Barrett adds that, based on conversations with Adaptive's clients, the most desirable outcome for both sell- and buy-side users is ease of application and workflow enablement. Large buy-side firms want to be able to breach silos—for example, allowing cost-effective, multi-asset execution—while the sell side wants to be able to supply this kind of execution functionality to their buy-side clients.

“Gone are the days of the full-fat, department store-style single-dealer platform being built by a universal Tier 1,” Barrett says. “I mean, they have them still, but I'm not sure that if you didn't have one, you would build a new one. Instead, they are now looking to deploy smaller pieces of functionality that are sticky with clients' workflows. So rather than having to build the entire department store, you build just the bit that really differentiates you as a business. FDC3 is about the interop, but the broader demand is about capability deployment on the user desktop.” 

State Street and FactSet look to take interop to the next level

Firms aim to provide a front-to-back data flow for asset managers, forming ‘deeper relationships’ with their clients. By [Wei-Shen Wong](#)

Achieving front-to-back interoperability between systems and platforms can be a challenging task. But as interoperability becomes more of a buy-side necessity, vendor firms are turning to each other to provide a more consolidated experience for their clients.

One example is State Street and FactSet’s most recent partnership. The two companies have jointly announced they are integrating State Street’s front-to-back platform Alpha with FactSet to streamline data flows across the front, middle and back office through State Street’s Alpha Data Platform (ADP).

Christopher Ellis, global head of strategic initiatives at FactSet, who has been with the company since 1994, tells *WatersTechnology* that clients are demanding that they have a “deeper relationship” with fewer partners and that those remaining partners work together.

So about a year ago, FactSet and State Street set out to understand how they could collaborate to essentially make life easier for their common clients.

“We met with 15 key clients globally last summer and asked them what they would want State Street and FactSet to do to help and be consistent with that theme,” Ellis says.

These clients were FactSet and State Street’s largest asset owners and asset management firms. According to Ellis, the result of those meetings came down to four pillars, which stretch the scope of what a typical partnership looks like. “I’ve never worked on a partnership that is this broad in scope,” he says.

Two of these pillars are primarily data coming from the State Street Alpha ecosystem into FactSet, and the remain-



ing two are data coming from FactSet into State Street. So, for example, State Street portfolio accounting data will be integrated into FactSet’s performance attribution, risk, and reporting capabilities, while portfolio analytics from FactSet will be integrated into State Street’s ADP. Additionally, portfolios in FactSet will update intra-day as trades are executed and allocated in Charles River, while proposed trades in FactSet transfer to Charles River for compliance and order management.

A pain point that asset managers grapple with is the lack of integration across their back-, middle-, and front-office operations, says Benjamin Quinlan, CEO and managing partner at Hong Kong-based strategic consulting firm Quinlan and Associates.

In an increasingly data-driven landscape, dealing with siloed data across different business areas makes it harder for firms to derive meaningful insights.

“The streamlining of data flows can help create a dependable source of ground truth for the buy side to leverage in a bid to gain a deeper understanding of their customers, enabling asset managers to turn those added insights into a more enriching customer experience to win over a greater proportion of heart share and wallet share,” he says.

The partnership between State Street and FactSet, Quinlan says, allows State Street to provide its clients with access to aggregated data, analytics, and real-time insights. “This not only translates to a better customer experience, giving State Street a competitive edge over its peers, but also allows the asset manager to collect data on how customers are interacting with its offerings. That treasure trove of data can create a positive feedback loop of incremental improvements,” he says.

Come together

The core of the partnership is ADP, which acts as the integration point between FactSet and State Street, says Bruce Feibel, platform manager for State Street’s Alpha platform.

“What we’ve done with FactSet at the bottom level of this is align our data model. So the data coming out of State Street’s portfolio, accounting, and valuation services is aligned to what FactSet needs to ingest in order to help our clients construct and analyze portfolios, and also get the data back into the ADP for integration with other datasets that are required to report to their end clients,” Feibel says.

Doing this manually creates challenges for asset managers, he adds.

“An example of that is performance attribution. It’s a pretty heavy lift to feed that process with the right data so that investment professionals can understand their portfolios with accurate data. State Street Alpha is the provider of the portfolio information—what do we own, what’s it worth, and so on. Aligning those data models so that these capabilities work together is what this

partnership is about; we want to analyze portfolios using FactSet's capabilities, but provide the integrated data from the ADP," Feibel says.

FactSet's standard datafeeds—including fundamental data, ESG data, estimate data, and so on—can flow into State Street's ecosystem for the user to use anywhere they choose.

"The beauty of ADP is that we're not connecting to [State Street's] TruView accounting platform, and we're not connecting to Charles River. We're connecting to ADP, and the State Street folks have built all of the interconnectivity from TruView to ADP and from CRD to ADP, so FactSet is just interacting with ADP," says Ellis.

The Snowflake connection

The ADP was built on Snowflake and Microsoft Azure, while FactSet has its own integrations with Snowflake. Ellis says this allows FactSet to push data into ADP through a Snowflake pipe, and then State Street can take that data and push it anywhere clients need it on the State Street ecosystem.

Attribution results, risk analytics, portfolio characteristics, complex weights composition numbers that are core to FactSet's performance, risk, and reporting capabilities, have a similar pipe into ADP.

Ellis notes that it's the clients' data, and not FactSet's data, that is going through those pipes. "We're making it easy for them to distribute and push it anywhere they want," says Ellis.

Feibel adds that while Snowflake certainly makes it easier than in the past to do these kinds of data exchanges by aligning State Street and FactSet's data models and their systems can communicate together so as to make updates in real-time, "that's a unique capability to State Street Alpha and FactSet working together that's very difficult for any client we have in common to do on their own."

Front-to-back and bottom-up

State Street has been at work on its front-to-back platform for a while now. With its integration with Charles River

Development, State Street can streamline workflows from where trades are made to trade confirmation, settlement, and security servicing, and then provide information back to the front office for trading, portfolio management, client reporting, and other functions.

State Street isn't alone in seeking that front-to-back play. Last year, Broadridge acquired Itiviti to help firms simplify their front-to-back technology stack.

Cubillas Ding, research director for capital markets and investment management at research and advisory firm Celent, says the front-to-back streamlining trends in the past few years focused mainly on the "portfolio implementation" aspects of investment operations, such as order generation, trade execution, margining, risk management, accounting, and other post-trade functions to achieve transaction settlement.

However, that is changing, he says.

"With initiatives like FactSet and State Street Alpha, you are now seeing spheres of operational and data convergence expanded to ensure that the 'thinking' aspects of the investment lifecycle—portfolio design and blueprinting activities, which include investment research, ideation, goal planning/target setting, asset allocation, strategy formulation, and top-down portfolio construction—are further aligned with the bottom-up implementation activities when orders/trades are being executed," he says.

This, in turn, will help asset managers achieve differentiated benefits, such as faster onboarding of new datasets, more compressed time-to-decision insights, and more dynamic portfolio design capabilities.

Achieving true front-to-back integration and interoperability, though, is easier said than done. Feibel says it can be challenging to standardize and streamline workflows across disparate activities. "The typical client that we've been able to serve with the platform is looking to rationalize, consolidate, and to create a platform that has secured their own growth, whether that's adding new asset classes, or incorporating acquisitions," he says.

To that point, Ellis adds that FactSet and State Street have connected their technical teams to build the plumbing between the two systems. He acknowledges that no matter how perfect the plumbing, there will still be some leakages. So, what they've done is taken the teams that handle the data scrubbing and the data follow-through and connected those teams and worked on the cadence of how they're going to work together so any unforeseen issues will not fall back on the client.

Interoperability is one of the fundamental tenets of the ADP, Feibel adds. He says, in the envisioning of State Street Alpha, the firm knew Alpha would only succeed if it could incorporate various sources of data and capabilities that firms like FactSet provide to State Street's clients.

"The platform is only valuable to the extent that we're open to being able to bring the data in and to feed those capabilities as they need to be fed so that our clients can use them," he says.

This continued push toward interoperability and data integration may lead to a mindset shift on the part of vendors, from breadth- to depth-focused, that is, developing deeper relationships with a smaller set of clients, as opposed to shallow ones with a longer client list, says Quinlan.

"To that end, vendors will have to show that they have a deep enough understanding of not just their client's needs and pain points in the present, but also of the direction that the client is aspiring to, and consequently that the vendor can continue to support them on their quest in the long run. As such, we may see a shift from short-term transactional thinking to more a long-term journey-oriented stance being adopted by vendors, as they look to form stickier relationships with an elongated lifecycle," says Quinlan. "Those vendors that are successful in forming these deep relationships can emerge as the big winners from the consolidation that may play out." **WT**

SocGen to move **datacenter footprint** in Americas to AWS, Azure

SocGen Americas plans to significantly reduce and even close datacenters in the US as a result of moving to the cloud and defining controls around cloud and data governance. By [Nyela Graham](#)



In the capital markets, the shift of compute power and data storage from on-premises equipment or third-party datacenters to the cloud continues full steam ahead. However, establishing good data governance in cloud environments is emerging as a challenge to making the most of cloud's performance-enhancing and cost-reducing potential.

Examples of recent cloud migrations among large banks and asset managers include: Fidelity's asset management arm, which has migrated 98.8% of its applications exclusively to Amazon Web Services (AWS) and plans to stop using physical datacenters in the near future; UBS, which has moved 33% of its workloads to the cloud; RBC, which is deploying a private-public hybrid cloud model and—as of earlier this year—had already migrated 600 applications, with 80% of them using

a private architecture; and JP Morgan, which expects to move between 30% and 50% of the bank's applications and data to the cloud.

Now, Societe Generale is joining that list. Simon Letort, chief digital officer and head of innovation for the Americas at SocGen, says the French bank aims to close its datacenters in the US “over the next few years” as it continues to migrate applications and its internal data lake to the cloud.

“Having most of the US-based applications on cloud will allow us to close down or significantly reduce our US datacenter footprint and no longer have to manage our own datacenters,” Letort says.

SocGen Americas started its cloud journey in 2015 with a gradual process that centered around examining how long data stayed on cloud. The bank started with capabilities like the ability

to “burst” to cloud, which the firm utilized to perform overnight batch processing for risk and profit-and-loss (P&L) calculations. Cloud bursting allows applications that typically run on a private cloud or datacenter to “burst” onto a public cloud when their computing capacity requirements spike. Once volumes stabilize, the data processing moves back to the private cloud.

The second step was moving pricing engines, such as for credit valuation, which allowed traders to launch computations on demand in the cloud.

“The third step is when you start to leave data on public cloud for a long period of time,” Letort says. “We started to do that when we migrated our client portal for research and analytics.”

The types of data stored on the portal include market parameters, market data, non-client data and non-

trading related data—in other words, no highly confidential data.

The fourth step involved moving into trading-related use cases, where trading applications would always be called on for high-reliability and trading production.

“It convinced us that this was a very beneficial investment, but it also showed us that this couldn’t scale,” Letort says. “We can’t wait months to approve an application that takes weeks to code.”

As a result, the bank had to look at new ways to improve the approval process. To address these inefficiencies, the bank set up a dedicated team to oversee “cloud control.”

“We want to get to a point where approving an app running on cloud or approving an app on-premises would be similar, and follow the same software lifecycle rules,” Letort says.

SocGen Americas utilizes both Amazon Web Services and Microsoft Azure clouds, but it took those services and “hardened” them for use, meaning the bank would build additional features on top of those in-house—for example, encryption tools around those services to ensure a higher degree of security. After about three years of development, the cloud control team has made 50 services available for faster approval of applications to the cloud.

Lessons learned

One component of SG Americas’ overall migration was moving its data lake to the cloud, prompted by frustration with issues around processing and limitations created by quotas. Letort described the key challenges of migrating the data lake to the cloud at *WatersTechnology’s* North American Financial Information Summit in May.

“Data scientists were running tick-by-tick analysis that would take, like, eight hours,” he said at the event. More nodes could tamp down the run time, but that could also mean the bank needed to add more servers, which would make operations more costly. Instead, Societe Generale spun up a

Google Cloud sandbox and supplied it with dummy data. Letort’s team used Google’s BigQuery data warehouse to run the test. The service allows analysis on petabytes of data and mirrors what the big data team at SocGen Americas was used to doing on-premises.

“In a few minutes, your cluster is spun up,” Letort said. The move to cloud was clear and the bank moved to shut down its on-prem data lake and move its wholesale data lake to Azure. The data lake had previously been housed on-prem in a datacenter in Paris.

“If we had decided to run the data lake on-premises at SocGen, we would have had to duplicate the infrastructure and staff required to run big data both in the US and in Europe, which is expensive,” he said at the event in May.

Today, he says any further cloud projects will be governed by the lessons learned during its migration effort so far.

“We decided as a firm to take a careful approach in reviewing all the datasets we will be transferring from on-prem to cloud and this showed that we kind of lacked standards,” he says, adding that establishing data governance rules is the “last mile” to a successful migration. “We’re still in an environment where it does take a certain amount of manual review, human review, legal, compliance before being able to send out the data and ideally to fully benefit from the power of cloud, you will have that streamlined and kind of automated based on some clear rules.”

An industry problem needing industry solutions

To help financial firms along their journey to cloud adoption, data management trade association the EDM Council last year launched its Cloud Data Management Capabilities (CDMC), an assessment and certification framework that aims to develop and implement standards and best practices for handling sensitive data within cloud environments. The framework includes six core components—govern-

ance and accountability; cataloging and classification; accessibility and usage; protection and privacy; data lifecycle; and data and technical architecture—along with 37 capabilities and 14 key controls for managing sensitive data in the cloud. Societe Generale was among the banks that contributed to the development of the framework.

However, Letort says that while the framework is helpful, applying it at scale to petabytes of data, either structured or unstructured, can be difficult when the datasets aren’t cataloged or classified, to begin with. “We lack a decent catalog of datasets, and associated classification,” he says. “I would be surprised to see many firms having a very good detailed, granular inventory of all their tables, columns, and rows of data.”

One reason banks may not have a cohesive classification and inventory of their data is that data within a bank can often be siloed, reflecting the way an organization has grown over time, says Virginie O’Shea, founder of Firebrand Research. “You’ve got banks that have been merged from lots of different entities, and you’ve got different desks that operate completely differently from each other. Fixed income and equities may not be using the same terminology, and there may be different taxonomies underlying their data because they were separate businesses,” she says.

Connecting these separate business areas and their datasets into a cohesive whole could be solved by manual development, though Letort says this would be expensive and is not scalable long term. Instead, the industry needs to develop tooling for data discovery and data classification, he says, citing Goldman Sachs’s open-sourced data management platform Legend as a tool that can work alongside the CDMC framework.

“I think this is the kind of approach that is needed in complement to the CDMC—and when you combine the two, you have something powerful,” Letort says. [wt](#)

Cboe migrations give exchange a foothold in Canadian equities

The exchange group is looking northward for its global expansion drive with the completed migrations of three trading system acquisitions. By [Joanna Wright](#)

Cboe Global Markets has made no secret of its plans for the domination of electronic derivatives trading—the exchange operator and market infrastructure provider is clear about its ambitions to provide the world's biggest derivatives and securities trading platform. The group last year launched a derivatives platform in Europe and is pushing ahead with plans to gain ground in Asia-Pacific and Canada.

Cboe has a playbook for moving into new geographies, what COO Chris Isaacson calls a “flywheel.” First, it gets a toehold in spot markets like cash equities, which produce market data that can form the basis of future indexes. Then, the exchange looks to acquire a clearing business to help it launch a derivatives platform: While US options are cleared on a utility (the Options Clearing Corp.), in most other listed derivatives markets, contracts trade and clear on platforms owned by the same group. Once the clearing services are in place, Cboe can launch its derivatives venue.

Though it took the better part of a decade, Cboe's expansionist flywheel has rolled a full 360 degrees in Europe with the 2021 launch of CDEX, a service for the trading and clearing of futures and options that runs alongside Cboe's existing Amsterdam exchange. CDEX leveraged Cboe's established cash equities business in the bloc, which provided the group with market data with which it could develop proprietary European indexes. When the time came to build the derivatives platform, this suite of indexes could form the basis for equity index contracts. Cboe's acquisition of cash equities clearing house EuroCCP gave it the clearing service it needed to offer CDEX.



“We recently launched European derivatives, so we've completed the flywheel. It's taken seven to 10 years to achieve that, but we have it now,” Isaacson says. “In the US, obviously, we have cash, data, and derivatives—we use the OCC for the clearing of derivatives. So that is our playbook, and we are forthright with investors that this is our strategy around the world.”

In Asia-Pacific, Cboe last year bought Chi-X, gaining access to the securities markets of Japan and Australia, its first significant foray into that region.

And in 2020, Cboe began looking across Lake Superior to the US's northern neighbor. Canada's attractions were obvious, Isaacson says, because it's open to competition, it's close to the US, and it has a similar regulatory regime. Many companies are cross-listed on US and Canadian venues.

“We look at the world and ask, where is competition allowed in equity markets? We didn't have a footprint in Canada, so in early 2020, we decided we wanted it to be part of our global expansion,” he says.

Making matches

Cboe's first foray into Canada was its purchase of MatchNow, the country's biggest dark pool. In 2020, Cboe bought the alternative trading system (ATS) from Virtu Financial, which itself had inherited it from ITG during its acquisition of the agency broker in 2019—MatchNow was part of ITG's dark trading suite Posit. At the time of the Cboe acquisition, MatchNow accounted for around 65% of Canada's dark trading volume, and about 7% of overall equity volume.

Cboe spent about a year and a half updating MatchNow's legacy tech and migrating it to its own tech and datacenters, announcing that the migration had been completed on February 1 of this year. At the same time, it announced that users of MatchNow's conditionals orders book would have access to improved conditionals capabilities from block trading network Bids Technologies, which Cboe had acquired in 2021.

“We thought that as part of this platform migration of MatchNow, we should also introduce Bids into the Canadian market. Almost 40 buy-side firms have onboarded, and another 180 have expressed interest. We have 26 broker-dealers that are accessing Cboe Bids Canada via algos,” Isaacson says.

This acquisition of Bids, a large US broker-dealer and block trading venue operator, was another piece of the puzzle for Cboe's move into Canada. Cboe and Bids were known to each other through a 2016 partnership to build Cboe Large in Scale, an indication-of-interest platform, in Europe. The 2021 acquisition gave Cboe more reach into the off-exchange segment of the US equity markets.

Then in June of this year, Cboe completed the third acquisition that it hopes will secure its toehold in Canadian cash equities. The group announced that it had bought Neo Exchange, a registered Canadian stock exchange. Together, MatchNow and Neo give Cboe about 15.6% of equity market share in Canada.

“Neo has built a fully fledged stock exchange in Canada. But they do more than just trading—they have data and they have a nice listings business. So we feel like this acquisition was the final step of three. We have MatchNow, Bids, and a stock exchange to get it to scale. We think this is a great foundation to build upon,” Isaacson says.

MatchNow, Bids, and Neo form the foundation of what Cboe hopes will be a holistic equity trading platform, Isaacson says, with a full range of trading mechanisms, from lit to dark, large to small, low frequency to high frequency. “With Neo, we have a very good lit order book and great listings venue. With MatchNow, we have a dark book that offers a set of different execution mechanisms, from those that are much more frequent and smaller, to Bids, which is less frequent and of much larger size, including conditionals,” he says.

While the migration of MatchNow is complete, Neo has yet to be fully integrated. Isaacson says Cboe is focused on the tech integrations of Chi-X in Australia and Japan, and once those are complete next year, the exchange group can look to complete Neo’s migration.

Good conditions

Cboe also plans to expand Bids globally, launching in Australia in Q1 2023 and in Japan in Q4 2023. Dark pools in Asia-Pacific are finding that tech can offer the ability to mask pre-trade information. As a source at an institutional broker-dealer in the region told *Waters Technology* back in March, microstructure improvements such as conditional orders, where clients can apply an “If I could, I would” type of order, as well as the use of minimum quantities, help clients get “chunkier” executions.

Cboe Bids Canada brings a sponsored access model to MatchNow: Buy-side firms can get direct electronic access to the platform via a sponsor of their choice—one of the platform’s network of sell-side subscribers—in order to send conditionals to the ATS. Sell-side and buy-side users can place orders through Bids’ desktop model, Bids Trader, which is integrated with the user’s execution and order management systems.

The conditional orders, which are two-stage orders much like indications of interest, are sent to the system, where they sit uncommitted until the platform identifies contra-liquidity and invites the user to firm up the order.

“Conditionals allow a user to be in two places at once with the hope of sourcing a block-sized transaction on Cboe Bids Canada. But they’re not doubly exposed because the order must be firmed up with a two-stage commit,” Isaacson says. The order is then executed on-exchange.

MatchNow had some conditionals functionality prior to Cboe, Isaacson says, but integrating Bids enabled Cboe to expand that for the buy side with the sponsorship model. MatchNow already had 56 existing broker-dealer subscribers when Cboe bought it; 33 of these participated in its conditional book. Those 33 are now participating on Cboe Bids Canada post-migration.

Cboe Bids Canada has launched a new feature called Willing To Trade aimed at making conditional trading easier for the sell side. This mechanism links large orders resting in MatchNow’s regular book with the conditionals book, so clients can let Bids handle the firm-up process. Cboe says this enables the subscriber to participate in the conditionals book without having to do any dev work as the buy-side clients are already onboarded and trading on Bids.

Headwinds

A source with decades of experience in senior roles at Canadian exchanges says Cboe’s acquisition of Bids makes sense for the exchange: “Bids has had some

success in building decent market share and daily turnover. Cboe can standardize the technology and messaging, apply the capabilities they have, take a profitable business, and make it more profitable, and bring those capabilities to more markets.”

As for Neo, the source says, the listings part of the business would have been especially attractive to Cboe. They say Canada has a lot of listed companies due to its extractive industries—junior mining companies need funding, and this is not a space that private equity ventures into much. Companies listed on exchanges pay annual fees, and so for the investment bankers advising exchanges, listings are an attractive source of regular, reliable revenue, unlike trading fees.

They say, however, that when it comes to the wider strategy of establishing a derivatives platform, Cboe might encounter some difficulty in the future.

It makes sense that Cboe would want to establish such a platform—derivatives platforms are far more profitable than cash equities, partly because listed derivatives remain in those silos, trading and clearing with the same group, giving that group ownership of, and chance to charge for, the whole stack. And Canada may seem attractive to Cboe given the presence of huge institutional investors, such as Ontario Teachers’ Pension Plan with roughly \$185.2 billion in assets under management, which could be a market for index options.

But Canadian derivatives markets, which are relatively shallow, compete with cash equities markets that have evolved to be deep, liquid, and efficient, or they would have lost liquidity to New York. “Because the ETF market is there, you have all these other instruments, you can drop C\$100 million (\$74 million) into 60 index stocks and there isn’t even a ripple. And five banks would be standing by to backstop the order,” the source says. “The virtuous liquidity cycle has never taken off.” 

Additional reporting by Wei-Shen Wong

Firms challenging Cusip file new lawsuit alleging antitrust violations, breach of contract

A broker-dealer and two investment managers have filed a new, joint complaint against the quartet of companies associated with management of the Cusip numbering system. By [Rebecca Natale](#)

Three financial firms have joined forces to file an amended class-action lawsuit against Cusip Global Services (CGS), S&P Global, FactSet, and the American Bankers Association (ABA), alleging violations of US antitrust laws, breach of contract, and violations of New York and Connecticut unfair business practice statutes, after a federal judge in the Southern District of New York in July ordered two previously competing civil suits brought by the firms against Cusip be consolidated.

New York-based broker Dinosaur Financial, Connecticut-based Hildene Capital Management, and Swiss Life Investment Management are seeking a ruling on whether Cusip numbers—de facto serial numbers used to identify nearly all North American stocks and bonds—are copyrightable or not, challenging the ABA's assertion of copyright over the Cusip numbering system, which it has held for more than 50 years.

The new suit also alleges CGS, S&P, FactSet, and the ABA violated sections 1 and 2 of the Sherman Antitrust Act by boycotting Cusip users who refused to enter into licensing agreements with S&P and FactSet, and by using the standard-setting process—Cusip is the national standard for securities identification under the American National Standards Institute's Accredited Standards Committee X9—to squash would-be competitors to Cusip, such as Bloomberg's Financial Instrument Global Identifier (Figi).

However, the suit sets up a fallback, stating that if Cusip numbers are found to be copyrightable, then the foursome is at least in violation of the ABA's commitment to X9 that it would offer licenses to Cusip numbers on fair, reasonable

and non-discriminatory (Frاند) terms, including price, which the suit says was and is a condition of X9's designation of Cusip as the national standard.

The two initial suits were each filed in March, days after FactSet closed its acquisition of CGS from S&P Global for more than \$1.9 billion. S&P had operated the CGS on behalf of the ABA for 53 years until the European Commission stipulated that the data giant must divest the business unit as part of its ongoing merger with IHS Markit.

Since then, the defendants have rejected the claims in court, particularly on the challenge to Cusip numbers' copyrightability. The basis for the challenge comes from a precedent-setting case in US copyright law—Feist Publications Incorporated vs. Rural Telephone Service Company—in which the US Supreme Court ruled that facts, or information alone, cannot be copyrighted without a minimum of creativity. For example, telephone numbers are facts and therefore cannot be copyrighted.

Though the plaintiffs don't dispute that the ABA holds copyright over CGS's master database of 60 data elements relating to more than 26 million financial instruments, they are challenging whether the underlying individual reference numbers that make up the database are copyrightable.

In a letter from June, S&P's lawyers wrote that the CGS subscription agreement grants users the right to access, download and use any of the data originating from the Cusip database, not just individual Cusip identifiers, and that anyone, including financial firms, can obtain and use individual Cusip identifiers for free from government and public sources such as the websites of the

Securities and Exchange Commission and its Edgar database, the Municipal Securities Regulatory Board and its Emma database, and the security issuers themselves.

However, some industry practitioners disagree. The CEO of a data vendor that licenses Cusip numbers for its own users says that one reading the CGS license page, which states that “CGS license fees for access to and usage of CGS's commercial products or services (and the CGS data contained therein) apply to both direct and indirect end-user customers” gets a different impression. CGS's own website says the company's data may be publicly available in some offering documents and from other sources.

“The problem, and what I think a lot of people in the industry are really ticked off about, is who has to pay to use it?” the CEO says. “[For example], DTCC has to pay to use it. Then DTCC takes those things—the trades—and they will ship it to BlackRock. BlackRock has to pay when they receive the data from DTCC, and then they have to turn around and pay for it again if they take that data and push it to somebody else. ...And so Cusip basically has this monopoly that says anybody who sends and anybody who receives it needs to have a license.”

However, all those different affected parties may find some solace in the new lawsuit. As part of the amended complaint, the suit now represents all entities who pay license fees to use Cusip numbers in their businesses, rather than only financial institutions. The new class includes institutions, data vendors, fintech companies, and other potential competitors who could develop innovative products and services using the Cusip numbers. [WT](#)

MDSL cuts data vets, retools to focus on technical sales

The loss of experienced data specialists has shocked some industry execs, but the vendor says its plan to train a new cadre of technically focused staff will result in better service overall. By [Max Bowie](#)

Calero-MDSL has closed its NYC office and laid off several senior sales and relationship management staff as it implements a more technical sales model worldwide.

Sources say the market data and telecoms inventory and cost management software provider shed more than 20 staff across a range of roles, including developers, IT, service delivery and telecoms expense management relationship management, as well as its remaining two New York-based market data specialists, and another member of the data team in Hong Kong. These follow the exits at the end of last year of experienced data specialists Laura Foggini and Frank Colombo. It will now service market data clients from its other global locations, including its Rochester, NY, and Phoenix, Arizona, locations in the US, its three offices in the UK, and its Asia office in Hong Kong.

The staff affected are said to have been “blindsided” by the move, which was unexpected, considering the company signed up a record number of customers for its market data platform. “None of us were non-contributors to revenue. We’ve all been in the industry for many years, and we’ve all bent over backwards to serve clients and make them happy,” says one former employee. “I was as puzzled as anyone.”

The company confirms some departures but says the move is part of a longer-term strategy to provide even better service to clients and meet their evolving and growing demands via a team with broader technical skills, says Alistair Brooker, general manager of the market data business line at Calero-MDSL, based in the vendor’s Tunbridge Wells office.



“We unfortunately had some long-termers leave the organization, but have backfilled with more technical, data-savvy team members to better assist our customers with the huge amounts of integration and automation that they ask of our MDM system. We have this technical account manager model working very well to support our customer demands in EMEA and have been working to set up a similar model in the US also.”

Despite “record” client wins over the past 18 months, “the customer support model in our US team wasn’t as assured as I wanted it to be,” Brooker says. “There were improvements we’ve wanted to make for some time. This wasn’t a recent snap judgement.”

The number of staff affected may be a fraction of the company’s roughly 850 employees, resulting from MDSL’s 2017 merger with former competitor telecoms expense management software vendor Telesoft and a 2019 merger with Calero. In addition, on August 2 this year, the company acquired Iowa-based telecoms expense management provider Network Control and its staff for an undisclosed sum. Network Control brings more than 60 employees to the company.

As the market data side of the company’s overall business has been diluted by multiple telecoms-related deals, “the perception is that they’re focusing more on telecoms than market data,” says one source familiar with the situation.

Calero-MDSL insists it is committed to market data. “In terms of our market data business, I couldn’t be more excited, with a record number of new customers coming to us in the last 18 months—the most we’ve seen in 15 years. It confirms to me that we’re on the right track with our roadmap and customer solutions around the MDM product,” Brooker says. “New team members are getting skilled in data analysis/engineering, customer account management, participating in the FIA training programs run by FID for market data knowledge, and are learning from the experienced team members around them who have 10 to 15 years in the industry. This is the model of team member that we’re looking for to answer our customers’ requirements and expectations for a system that is now expected to handle more than just market data cost allocation.”

But end-users aren’t as excited by the news. One former MDSL client said they would be concerned by the loss of experienced staff. “Some of those people have been there forever. They know every inch of that software, and in some cases had set up special workflows specifically for me,” the source says. “That industry experience is a big deal.”

Nevertheless, Brooker insists that once the changes shake out, the vendor will have a better sales, account management and support model, and its clients will reap the rewards, noting that the company continues to grow and make additional hires. [wt](#)

NEWSDESK

WatersTechnology's roundup of headlines that hit the wire this quarter from around the industry

Numerix acquired by Genstar Capital



Numerix to expand post-acquisition

Numerix, a provider of a variety of risk management tools geared toward the capital markets, has been acquired by private equity firm Genstar Capital. Numerix has 19 offices, 225 clients and 90 partners globally.

Genstar invests in companies in the financial services, software industries, industrials and healthcare sectors.

"For over 20 years, we have worked tirelessly to provide the most sophisticated analytic and technology solutions available to capital market participants across the globe," said Steven O'Hanlon, CEO and president of Numerix, in a statement. "We've always had a clear vision for the organization. Together with Genstar and their long track record in the financial services and software industries, we look forward to expanding our footprint across the entire value chain in the front-to-risk market providing even greater value to our customers."

Scott Niehaus, director of Genstar, said: "Numerix is capitalizing on several macro tailwinds, including the digital transformation of capital markets, continued complex regulatory requirements and substantial market volatility across asset classes and geographies."

Terms of the deal were not disclosed.

Cboe to use Snowflake cloud for analytics



Cboe moves from on-premises to cloud

Cboe Global Markets will migrate its corporate data and analytics functions from its on-premises systems to the Snowflake cloud.

The pairing will allow the exchange to "accelerate our cloud migration journey and modernize how we manage and utilize our proprietary corporate data and analytics for our benefit and that of our customers," said

Eileen Smith, senior vice president for data and analytics at Cboe. "Snowflake offers a highly flexible and scalable data cloud platform that will enable our teams to handle increasing amounts of data with greater speed and efficiency, improving our capabilities for managing, mining, and analyzing that data internally across our global business and sharing those insights with our customers. We look forward to the many expected benefits that Snowflake will provide."

Broadway Technology integrates with MarketAxess



The joint feature is currently live

Broadway Technology, a front-office solutions provider, has partnered with MarketAxess. The electronic trading platform provider and Broadway have collaborated to build and support a US Treasury request-for-quote (RFQ) workflow feature in the MarketAxess US Treasuries marketplace. The feature is currently live and available to customers of both Broadway and MarketAxess. The new integration builds on the existing MarketAxess Live Markets Rates partnership and now includes RFQ, in aim to offer mutual dealer clients the ability to respond to inquiries from MarketAxess' customer base.

ING Ventures invests in OpenFin

OpenFin has secured an investment from ING Ventures. The venture capital arm of the bank intends to accelerate the expansion of OpenFin operating system throughout the financial industry. This announcement follows last year's launch of OpenFin Workspace, the visual interface of OpenFin OS, which includes components for complex windowing, advanced search, actionable notifications and application discovery.

SmartStream and Kynec collaborate to deliver integrated OTC margin solution

SmartStream has formed a strategic alliance with Kynec, a front-office clearing solutions provider. The transaction lifecycle management service provider has collaborated with Kynec to deliver consolidated margin positions across global over-the-counter (OTC) cleared and bilateral markets.


This alliance intends to integrate cleared margin data from Kynec's Rubicon platform with bilateral OTC margin data from SmartStream's Transaction Lifecycle Management (TLM) Collateral Management solution. Collateral operations can now access this information for workflow management and reporting, while the front office will be able to access information directly to manage liquidity, margin funding, and collateral optimization across bilateral and cleared margin positions.

DTCC's Project Ion platform goes live in parallel production environment

The Depository Trust & Clearing Corp.'s (DTCC) Project Ion platform has gone live in a parallel production environment while upholding the firm's safety standards. The platform, which is being developed as an alternative settlement platform leveraging distributed-ledger technology (DLT), is now parallel processing an average of over 100,000 bilateral equity transactions per day. The DTCC subsidiary, the Depository Trust Co.'s (DTC's), classic settlement systems remain the authoritative record.

OPEN OUTCRY

“



“Our data was becoming unmanaged and uncontrolled because we didn’t have a central repository. We noticed things like company records going stale, or missing data.”

Chris Caliri, RA Capital

» see page 40 for full feature...

“

“None of us were non-contributors to revenue. We’ve all been in the industry for many years, and we’ve all bent over backwards to serve clients and make them happy. I was as puzzled as anyone.”

Former MDSL employee

» see page 23 for full story...



“

“We want to make a bigger statement about how the core CDM is open-source, and other trade associations, should they want to develop things for different markets, can similarly build on this neutral CDM.”

Ian Sloyan, Isda



» see page 4 for full story...

“



“Sometimes in a legacy platform, the code has grown over 20 or 30 years, and nobody quite knows what it does. But when you implement a new platform, or a new component—like position management, for example—we can implement that as a cloud-enabled component. And it doesn’t have to be everything in one shot.”

James Marsden, Broadridge

» see page 44 for full feature...

“

“When you get into the reality vs. the plan, if you didn’t change anything, then you’re probably not right. ... As we’ve gone into it for real and spoken with their engineers in depth, we’ve learned a few things and changed our plans in some ways. That shouldn’t impact the timing, but will impact ultimately how and what we do, and when.”

Brad Levy, Symphony

» see page 48 for full feature...



“



“Traditional data initiatives tend to fail because they have similar characteristics and therefore often look the same—we call it the ‘big red blob.’ The red blob could be a warehouse, a data lake, a lake house, or a lake house on the cloud!

What I often tell people

when they’re looking at a traditional data initiative diagram is to lean back a little and squint your eyes: You’ll see the same red blob over and over.”

Rick Carey, UBS

» see page 32 for full feature...

“

“One of the things that would always frustrate me 10, 15, 20 years ago, was: Let’s say you were a fund that traded one asset class—so you’re an equities shop and you wanted to start trading fixed income. Oftentimes, to do that you would buy an entirely new vertical of systems and you’d have to put them together for that asset class.”

Neal Pawar, Qontigo

» see page 52 for full feature...



waterstechnology



View our latest content
even when offline



View articles offline, save and share favourite stories, and customise the type of content you would like to receive with our **Waters Tech Live app**.



waterstechnology.com/apps

Sell side ramps up outsourced trading desk services



A raft of new entrants are offering outsourced trading services to buy-side firms, anticipating a wave of takeup among larger asset managers, driven by cost and coverage needs. But are they aiming too high? By [Max Bowie](#)

In the fall of 2020, right as the second wave of the Covid-19 pandemic was hitting, Frontier Road opened its doors (virtually) in London.

As is true of any startup, the asset manager had to decide whether to buy or build. For Martin Bercetche, the choice was simple: He is running an investment firm that focuses on the niche sector of emerging markets credit; it's not a tech company, so why not leverage the expertise of an outsourced trading provider? That led the firm to New York-based Cowen Group.

"Financing Angolan government bonds isn't everyone's cup of tea," says Bercetche, Frontier Road's chief investment officer. "A long-short strategy

needs good financing and repo abilities. There aren't many specialists in emerging markets credit from an outsourced trading desk perspective. But Cowen has some experienced emerging markets credit traders who understand how the game is played. I'm confident I can get the best price in the shortest time."

Cowen Group has 45 traders on its outsourced buy-side desk, with connections to 152 brokers worldwide, across multiple asset classes. All Cowen's outsourced traders use a combination of SS&C's Eze OMS and RealTick EMS, and has built its own security master in SS&C's Advent Geneva back-office platform, in which it has also built workflow tools to generate and distribute

automated reports to clients. The firm also built a client-facing prime brokerage portal, and a "lite" version of its own P&L monitoring system for clients who don't want to build or buy their own.

While most adoption of outsourced trading has so far centered around equities and gaining access to international equities market, emerging markets credit is heavily reliant on expertise and manual processes. As such, Cowen has been investing heavily in its fixed-income capabilities—and human capital—over the last year, says Michael Rosen, managing director and global co-head of prime brokerage at Cowen Prime Services.

"The equities markets are very transparent and easy to navigate, but

fixed income is still more opaque and voice-brokered,” he says. “So, having a desk that’s connected to between 40 and 45 liquidity providers globally gives clients the reach that they couldn’t have achieved themselves. There’s a lot more value there than, for example, in the equities market.”

That expertise—coupled with the execution—has made getting Frontier Road off the ground much easier than it would otherwise have been. Bercetche notes that in emerging markets credit, you can’t simply call five banks and expect to get a price: “You have to find a real buyer.” He struggles to put a return on investment on the cost savings or performance gained, but the fact that Cowen has opened the firm up to 10 times the counterparties he would have had access to on his own has helped to jumpstart the firm’s trading activities.

“For a startup fund in its first year of existence, outsourced trading is very attractive because it lets me gain access to lots of different players—and to their research and expertise—that I wouldn’t otherwise have access to,” he says. “I wear a lot of hats, so this saves me a lot of time. Usually, I use Bloomberg chat to tell them the bond, size, price band, and what I want to do, and they go and work the order.”

Another bonus feature is that Cowen also caters to Frontier Road’s financing and regulatory operations needs. “So, it’s a one-stop shop,” Bercetche says.

Slow burn

The drumbeat for outsourced trading has been growing for several years now—and that only got louder during the pandemic. But while there’s a lot of talk about outsourced trading, only a small proportion of the buy side appears to be taking advantage of those benefits so far.

“Outsourced trading is a topic that gets a lot of attention and comes up pretty frequently: Both buy-side firms and outsourced trading desk providers call us up and ask what we see their peers doing. But, based on the results of our studies, it seems like outsourced trading is still quite nascent,” says Brad Tingley, research manager in Coalition Greenwich’s market structure and technology practice.



According to Coalition Greenwich’s *2020 Market Structure & Trading Technology Study*, only 10% of buy-side firms already outsourced their trading desks—specifically, that figure was for international trading, one of the most compelling use cases. For other areas and asset classes, that number was even lower.

And Greenwich’s latest report on buy-side spending, released in Q2 this year, reveals that the budgets for trading desks on the buy side increased by 5% overall in the last year, with the average firm spending \$2 million per year on technology and staff, making for \$10 billion per year spent on this area by the buy side overall.

But the research also shows that the increased spend was mostly driven by fixed-income desks, where budgets increased by 12%, whereas equities desks increased by 4%, and budgets assigned to currency trading desks decreased by 2%. If one assumes that firms spend more on the areas that make money and less on areas where they perceive less advantage, then these figures might provide a hint about which areas are more or less likely to be kept in-house or opened up to outsourcing arrangements.

“For most buy-side firms, their FX trading is just hedging their currency risk. It’s not a source of alpha, so they’re probably happy to outsource that execution,” Tingley says. “As it becomes more time-tested and outsourced trading desks can show a longer track record and that

they help firms’ investment management requirements, and as more buy-side firms start using it and report positive experiences, then we’ll see it expand into other areas.”

Any evidence that outsourced trading desks are the way of the future for the buy side is largely anecdotal. Today, the sector feels a bit like the movie *Field of Dreams*—if you build it, portfolio managers will come. It’s this thinking that has led to something of an arms race among established outsourced trading providers and newcomers, alike.

The pitch in favor of outsourced trading largely revolves around cost savings, fewer (manual/human) errors and streamlined workflows. Shifting that cost burden onto third parties that already have infrastructure and licenses in place is one way to take costs out of the industry overall, says George Black, a partner at Capco and US head of the consulting firm’s capital markets domain. In his view, the rationale for outsourcing at any firm should be that if something is not differentiating your business, you should be willing to pay someone else to do it for you.

“In the past, traders were seen as a value-add. But as liquidity and electrification in the equities markets have increased, and spreads and commissions have decreased, it has begun to be less of a differentiator, so trading has become more about expense management,”

Black says. “You have portfolio managers who decide what a portfolio should consist of, and the traders just act on those instructions. So firms are asking whether they should be investing in those roles.”

The mere suggestion that any portion of the buy side is ready to shift to using outsourced trading desks is creating intense competition among those offering outsourced trading services. These include custodians such as Northern Trust, specialists such as Tourmaline, firms like Cowen Group, large banks like UBS, and startups like Meraki Global Advisors.

The analytics angle

Proof that there’s a tipping point emerging in the outsourced trading space perhaps can be seen by the entrance of UBS. The Swiss bank had been providing a similar offering to around 300 sell-side clients across the US, Europe and Asia for years, but at the start of 2021 the bank decided to expand its service to the buy side.

By December 2021, it had constructed a new trading portal, covering equities markets. UBS will add fixed-income coverage at the end of June and plans to add listed derivatives capabilities in Q4 of this year.

“There’s a conversation on the buy side about what is your core business, and what core capabilities do you want to keep in-house vs. what can you benefit from outsourcing? For example, if you mostly trade in the US, but also want to trade in Asia, you can meet that capability,” says Mark Goodman, head of UBS Execution Hub and platforms in London. “By using an outsourced desk, you get exposure to that expertise.”

For its trading platform, UBS partnered with FlexTrade and with Virtu’s ITG division for the broker’s transaction cost analysis capabilities, both of which are already well used by buy-side firms. By buying systems off the shelf and customizing them rather than building something from scratch in-house, UBS was able to move quickly to take advantage of the growing opportunity, while offering broader capabilities to clients. For example, it is using FlexTrade’s API to plug machine-learning models, external data sources and capabilities into the platform.

“The API was so important because we wanted to take algorithms and apply them to the outsourced trading desk,” Goodman says. “Our aim in using the flexible API was to build out our machine-learning capabilities to optimize trading. So, the buy side doesn’t just have access to experienced traders, but we also have a very large dataset, which we can use to optimize a trade. For example, if, on a specific Friday in August, you want to trade a small-cap Swedish stock, how is best to execute that trade, and with which broker?”

Goodman believes UBS will entice the buy side by incorporating its data analytics capabilities to drive insights to the buy side. By running machine learning—rather than standard statistical analysis—against the dataset, which comprises all orders placed and trades executed by UBS, as well as data from

“The equities markets are very transparent and easy to navigate, but fixed income is still more opaque and voice-brokered. So, having a desk that’s connected to between 40 and 45 liquidity providers globally gives clients the reach that they couldn’t have achieved themselves. There’s a lot more value there than, for example, in the equities market.” **Michael Rosen, Cowen Prime Services**

clients, the bank can spot patterns and make sense of the large datasets.

“We’re in the early stages, and are starting to experiment with how to make it useful,” Goodman says. “But as the dataset gets bigger, we expect the quality will improve.”

UBS has staff serving this specific space in New York, the UK and Singapore, and will have a dedicated team in Frankfurt up and running in July. “We wanted a single, global order management system so we could provide a follow-the-sun service across our global team,” he adds.

Growth ‘supplements’

That idea of global coverage is one of the key drivers of interest in outsourced trading desks among buy-side firms, most of which tend to trade in their domestic market. But as clients’ portfolios become increasingly diverse and international, those managers may need to execute

trades in asset classes where they have limited expertise, and in regions or on markets where they don’t have a physical presence or a trading membership.

For Atla Capital Management, a San Juan, Puerto Rico-based startup investment fund focused on listed emerging market real-estate investments, the global coverage of an outsourcer—in this case, new entrant Meraki Global Advisors—to provide access to far-flung emerging markets “where trends around urbanization are most pronounced” was a no-brainer, says John Haskell, chief investment officer at Atla.

“I came to the decision to outsource my trading very early. The resources required to spin up a trading desk with the same quality and reliability were beyond my scale. And I didn’t want to bootstrap it internally because it’s an important function, and I don’t have

time to do it myself—I have many other demands on my time. It’s a question of what to do yourself and what to outsource—and outsourced trading was one of the more obvious decisions I had to make because the economics were obvious,” Haskell says.

Haskell has used Meraki since the inception of the fund, and describes the workings as being just like having an internal desk. His resume includes stints at larger buy-side firms that had internal trading desks, where he communicated with traders exclusively by phone or email. The relationship with Meraki is no different, he says.

Meraki was founded three years ago. The firm has offices and 16 clients spread across North America, Asia, and Europe, and is adding new clients at a rate of one every two months, says the company’s COO Michael Ashby. Meraki has found a sweet spot provid-

ing multi-asset outsourced trading for hedge funds with between \$1 billion and \$5 billion under management.

What makes Meraki somewhat unique, Ashby says, is that the firm sets up a desk for each client, and its traders handle a maximum of three to five clients. They don't share any resources, and—unlike other firms that can use their scale to gain an advantage with the broker-dealers with whom they trade—the sell side sees exactly which fund they're trading with, getting them a level of recognition with those firms that has declined under existing agency-broker models, he says.

A contributing factor to Meraki's growth during the Covid-19 pandemic is that new funds found it hard to raise funding over this time, which substantially increased the importance of firms like Meraki that can help facilitate capital introductions, Ashby adds.

Muddy waters

Rob McGrath has spent three decades in senior trading roles at buy-side firms, including as global head of trading for Abu Dhabi Investment Authority and as global head of trading for Schroders Investment Management, where he led a team of 42 traders with over \$500 billion in assets. He was also head trader at Oppenheimer & Co.

For much of that time on the buy side he believed trading functions didn't add value and should be outsourced. The problem, he says, is that age-old distrust between the buy side and brokers.

"For 15 to 20 years, I've firmly believed [outsourcing] is the way to go, but I don't think you can solve this by just having an asset manager give a trader an order," says McGrath, who today runs fintech consultancy ZigIQ. "When I was on the asset management side, I would give a trader an order to execute, but I wouldn't want to give them too much information that would give away my strategy—and yet I would still expect them to do a good job. How can they do that? The problem is, there's one piece of data at the asset manager, there's another piece that they share with the broker, then there's the data that the broker has—how can you expect to get the best outcome if you don't share data?"



Getting the best result, he says, requires access to all the data that an asset manager has—specifically, their models, strategies, portfolio data, and so on—rather than a single order, given with no context around it.

"I don't see buy-side firms giving that to a broker. But they might give it to a technology provider," he says.

That's where McGrath sees the next opportunity: to build a solution—initially targeting buy-side firms outsourcing their execution, but with the potential to expand to serve other market participants also—that can capture data from all sides and use it to set strategies, then monitor and adjust them in real time, in response to changes in the data. He's currently talking to potential partners, and hopes to have a product available in the fourth quarter of this year.

There's certainly room for more players on the outsourced trading court. Cowen says its own business in this space has doubled over the past three years, and expects to continue experiencing double-digit growth. That may not be the same for everyone.

"Whenever there's a 'shiny new toy,' you get a lot of firms hanging out a shingle," says Cowen's Rosen. "There might be brokers with excess capacity who think by saying they're in outsourced trading that they'll be a miraculous success. So, I think there will be a wave of new entrants—but I also think there will be a number of those that aren't able to meet clients' needs."

Looking to the future, the distinct forces driving different types of firms toward outsourcing their trading desks create interesting decisions for firms in a few years' time: Those choosing it to save money over existing resources will have to decide whether outsourcing has met their cost and performance expectations, and whether to continue or deepen those relationships. Meanwhile, those using it to gain broader reach during their startup phase will have to decide whether there's any advantage to bringing the function in-house—and whether they could replicate the same levels of service, coverage and reliability as those offered by an outsourcer.

Those drivers—and that decision—may depend on the size of the firm in question and what they're trying to achieve. And while outsourcers claim to be attracting larger clients, smaller and startup funds seem to stand to gain the most benefits.

"In the asset management industry, those with scale have a natural advantage. But there are countervailing trends, such as outsourcing, that give smaller firms a shot at success," says Atla's Haskell. "Those who break down the barriers to entry will help breathe life into the industry and drive innovation."

When it comes to outsourced trading, size matters

Size and scale are the key arguments on both sides of any outsourced trading desk agreement. The bigger an outsourcer, the more services it can offer. The more valuable a partner it is to buy-side firms, the more it can earn from offering those services. On the other side of the fence, potential buy-side clients are looking to increase their scale without increasing their size and expenses.

So, it's not surprising that larger firms such as UBS are getting into this space. This allows them to offer more banking and custodial services, in addition to those such as Northern Trust, which has an established offering for buy-side clients. Thus, buy-side clients can offload more aspects of their operations that are non-core and don't contribute to strategy and returns, while sell-side firms can cross-sell more services into which they already have sunk investments.

Northern Trust offers a number of services as standard that outsourced trading clients can utilize, including trade settlement and Swift messaging, calculating and executing currency legs of trades, reporting and commission management, and transaction cost analysis, among other services.

“We had a set of tools and services ... but if you integrate those into a client’s trading operations, you can raise the value of the proposition,” says Grant Johnsey, head of integrated trading solutions for the Americas at Northern Trust. This can include outsourcing

the sell side for resources such as to supplement their own in-house investment research, why not also leverage third parties to improve access to liquidity and execution?

Why not, indeed, says Jeff Estella, principal of Boston-based consulting firm Estella LLC, who previously spent almost 30 years on the buy side, and who has advised companies on this transition.

“Most large asset managers already have custodial middle- and back-office outsourcing relationships. The question is whether the desire exists to move into a front-office outsourcing relationships

ent firm, or broker-dealer,” he says. “In a supplemental world, you can have 100% of market hours covered, and have eyes and ears focused on orders during live market hours.”

As well as delivering broader coverage, reach and expertise, “it definitely improves clients’ costs, because you’re taking a fixed cost and turning it into a variable cost,” says Cowen’s Rosen.

For example, George Black, a partner at Capco and US head of the consulting firm’s capital markets domain describes buy-side clients that were able to save around \$400,000 per trader because their execution activity didn’t justify that level of costs. But there are certainly other advantages to be gleaned: Equally importantly, these firms gained increased levels of control over their operations and other costs overall as a result of needing to impose more structure and order for the outsourcing relationship to integrate seamlessly.

“When you outsource, you’re forced to implement more structure, and to think about things—for example, handoffs—in a different way,” Black says. “The overall results can be reduced numbers of failed trades, reduced capital costs, reduced incorrect allocations or settlement instructions, and a reduction in manual, human errors.”

One of the most-reported sources of errors among buy-side clients is foreign exchange (FX) components of trades, where firms need to calculate conversions and hedges for each account, says Northern Trust’s Johnsey, whereas an outsourcer can perform currency calculations faster and assume that currency risk.

This could be the starting point for larger outsourcing deals that the major banks and brokers envisage, says Coalition Greenwich’s Tingley.

“For most buy-side firms, their FX trading is just hedging their currency risk. It’s not a source of alpha, so they’re probably happy to outsource that execution,” he says. “As it becomes more time-tested and outsourced trading desks can show a longer track record and that they help firms’ investment management requirements, and as more buy-side firms start using it and report positive experiences, then we’ll see it expand into other areas.” **wt**



“There’s a conversation on the buy side about what is your core business, and what core capabilities do you want to keep in-house vs. what can you benefit from outsourcing? For example, if you mostly trade in the US, but also want to trade in Asia, you can meet that capability. By using an outsourced desk, you get exposure to that expertise.” Mark Goodman, UBS Execution Hub

back-office functions as well as front-office trading, and can also include supporting a firm’s alpha generation process via Northern Trust’s Equity Data Science (EDS) platform, he adds.

“EDS takes any data source the investment team is using, pulls it into a single database, then allows them to integrate their own price targets, take the data they’re analyzing, and weight it to their investment process,” he says. “It takes all the data into a common platform ... that allows them to codify their investment process. So, we’re not telling them what stocks to pick, but it takes research management, idea generation, and gives feedback to the portfolio manager based on their own data.”

Cutting out unnecessary overheads can be an important strategy. With pressure on margins and commissions down, outsourcing is becoming more widely accepted, says Tim O’Halloran, managing director at Tourmaline Partners, which provides outsourced trading services for buy-side firms. Since the buy side has always turned to

with a custodian,” Estella says. What should be driving that decision is an asset manager’s fiduciary responsibilities to clients.

Additionally, Tourmaline’s O’Halloran says that supplementing existing teams within larger buy-side firms that need to expand their teams urgently or that want access to a broader range of brokers is one of Tourmaline’s fastest-growing areas. “It’s a valuable resource to have at the ready during times of volatility,” he says. How valuable? Tourmaline traded just under half a trillion dollars in notional value of equities trades last year, and around \$40 million in derivatives contracts.

Estella likens this to the “elastic capacity” offered by cloud computing, which is already widely adopted across capital markets. The difference is that instead of dialing up or down virtual storage or processing power, you’re adding or reducing people power.

“There may be a region of the world—such as Asian equities, or Latin America—that an asset manager doesn’t cover as well as a custodian, independ-



Beyond the ‘big red blob’

UBS has started a journey to implementing mesh data architecture, a new concept for federated data availability in enterprises aimed at getting value out of analytical data. By **Joanna Wight**

If you work for a large financial firm, it's safe to assume that your data architecture is probably what UBS Group CTO Rick Carey calls a “big red blob.”

Like many technologists whose job it is to explain complex concepts to businesspeople, Carey employs some neat formulations, many of which, one senses, have attained the status of in-jokes at UBS.

The “big red blob” quip derives from illustrations in internal slide presentations in UBS's official colors—the Swiss bank's logo is black and red. Carey uses the phrase as shorthand for the kind of centralized, monolithic data architectures with which large enterprises have for decades tried to consolidate the data they take in from numerous sources.

These kinds of architectures will be familiar to *WatersTechnology* readers. Large organizations like banks and asset managers tend to employ various data warehouses and data lakes. Data warehouses are repositories of information standardized and managed for business purposes and fed by extract, transform and load (ETL) pipelines. The warehouse was succeeded by the concept of data lakes, a solution for the era of big data—vast pools of unstructured and semi-structured data, from which users can develop machine learning and AI models, if they can effectively sift through and evaluate this raw data.

Various combinations and hybridizations of these architectures have become buzzwords, especially as organizations are increasingly shifting their entire estates to the cloud. But experts, Carey among them, say these centralized, monolithic constructs can no longer deliver value to enterprises as they scale.

“Traditional data initiatives tend to fail because they have similar charac-

teristics and therefore often look the same—we call it the ‘big red blob.’ The red blob could be a warehouse, a data lake, a lake house, or a lake house on the cloud!” Carey tells *WatersTechnology*. “What I often tell people when they're looking at a traditional data initiative diagram is to lean back a little and squint your eyes: You'll see the same red blob over and over.”

“Traditional data initiatives tend to fail because they have similar characteristics and therefore often look the same—we call it the ‘big red blob.’ The red blob could be a warehouse, a data lake, a lake house, or a lake house on the cloud! What I often tell people when they're looking at a traditional data initiative diagram is to lean back a little and squint your eyes: You'll see the same red blob over and over.” Rick Carey, UBS

Organizations produce more and more data, and find more and more data from external sources to consume, especially as they grow and evolve. New tools have made it possible to capture more and more information for actionable insight and to improve customer experience.

“The size, sources, and types of data are growing every day. Take the movement of a mouse. It can be tracked to gauge the interest level of a person browsing a webpage: Do they hover? Do they move away? All this information gives us a great understanding of a customer's experience. The movement of a mouse—think of how much data that is!” Carey says.

Executives are dazzled by the possibilities of all this data, especially as the tools

for making use of it—analytics, storage capabilities, processing speed—get better and cheaper, and pour investment into tech and people. These executives like to talk a big game about shifting to a data-driven culture. But throughout financial services and beyond, no one seems to be succeeding with their big data platforms.

Consultancy New Vantage Partners published its 10th annual report on data and AI in January, for which it had surveyed 94 blue-chip corporates—including UBS, as well as JPMorgan Chase, Citi, Wells Fargo, and MetLife. The survey found that while investing in data and AI is growing, businesses are reporting failures on some crucial metrics of success. Of those canvassed, just 19.3% reported that they had established a data culture, 26.5% reported they have created a data-driven organization, and 47.4% reported they were competing on data and analytics.

So why are sophisticated organizations that are cognizant of the value of data and are pumping resources into storage and analytics and engineers failing to realize benefits?

This was a question that, in the 2010s, was troubling consultants at Thoughtworks, an influential software and delivery company known particularly for its central role in the development of the agile methodology.

In studying both successful and failed implementations, one of these consultants, Zhamak Dehghani, observed that enterprise data is siloed, with operational and analytical data in completely disparate “worlds” within the organization. Connecting these two worlds is an inefficient and fragile labyrinth of pipelines, underpinned by data warehouses fed by cumbersome ETL processes and housing data organized under universal, canonical models

too inelastic to respond to a dynamic organization, with its proliferation of applications and users, and data.

Dehghani knew that organizational and data architecture is inextricably linked, and she also observed that these data siloes were intermediated by teams of platform engineers who had to become hyper-specialized in the tools that serve big data applications—particularly as clients increasingly want to take these concepts to the cloud. The resulting talent shortage is not news to anyone in this industry, and it is a battle that no business will ever win: As the business grows, it must keep finding more engineers.

Also, according to Dehghani, data engineers are estranged from the producers and the users of the data, and so have little ownership of it; they barely understand how it was generated or the uses to which it will be put. This engineering layer creates a bottleneck for data to flow through the enterprise, and their solutions leave data consumers unsatisfied.

In 2018, Dehghani conceived a new approach, which she detailed later in a blog post. Dehghani wrote that the world needs to shift away from the paradigm of a data lake or a data warehouse to a distributed architecture that considers business domains as owners of data and treats data as a product. Dehghani, anxious to avoid the aquatic metaphors that characterize data management nomenclature, called her new methodology “data mesh.”

Data mesh as an approach is not a technology or a vendor solution or the result of a one-off consultation with a company like Thoughtworks. Rather, it is a tech-agnostic, federated architecture, built using a set of principles, that ensures the on-demand availability of data by users within a business on a peer-to-peer basis.

Thoughtworks did not respond to a request for an interview. But Dehghani has expanded her initial ideas into a book called *Data Mesh: Delivering data-driven value at scale*, and the first chapter is available for free on the Thoughtworks website.

It's early days for data mesh, but a small community of evangelists has sprung up, and companies including retailer

Zalando and trading platform provider CMC Markets have detailed their experiences of transitioning from the big red blob toward a data mesh architecture.

We mesh well

Data mesh certainly has an enthusiastic convert in Carey, who is excited about the possibilities of the approach.

“You have increasing sources of data and an increasing number of consumers—consumers almost always become producers [of data]. This tends to go beyond the limit of traditional data initiatives. And that's where mesh data architecture looks to be the best solution because ‘big red blobs’ struggle to keep up with the growth of data and the growth of the consumer,” Carey says.

For Carey, data mesh complements UBS's existing trajectory to becoming data-driven, which is accelerating under the leadership of CEO Ralph Hamers.

Hamers was brought on to UBS in 2020 as a digitizer. Under Hamers, UBS has started moving more employees to an agile way of working, extending the firm's construct of small, interdisciplinary teams called “hybrid pods” to more of the business, under a program known as Agile@UBS. Hamers said in the group's Q1 2022 earnings call that 10,000 staff had so far been shifted to agile under this program.

Late last year, the press reported on a memo in which Hamers told staff that UBS would be creating a new bank-wide team—called AI, Data and Analytics (ADA)—to manage data to use more AI and analytics, with the aim of attracting more wealthy clients in a highly competitive environment.

Carey says data mesh will inform the architecture of ADA.

“Data and consumers are increasing, and methodologies like AI, machine learning, analytics, computational capabilities, GPUs, and the cloud are also increasing. That to us is the sweet spot for a data mesh architecture,” Carey says. “Why do we think this? A data mesh architecture starts with the premise that data can be anywhere—and it already is everywhere.”

Carey uses the metaphor of a screen, such as you might find on a door to keep the bugs out. Or perhaps one could

think of it as a fishing net, made up of knots and strings that run between them.

“It's made of interlocking nodes, which have relationships with each other. In other words, the data is the ‘what,’ the consumer is the ‘who,’ the methodology is the ‘how,’ and the mesh contains the paths that bring all these together,” Carey says.


In a data mesh, the nodes are the products—data as a product is a fundamental principle of this approach. Data is a product not in that it's something ready for sale, but rather in the sense that the customer's needs come first in its design and packaging. If data is a product, Dehghani writes, it can be shared directly with users who are peers, such as analysts, avoiding siloing and disintermediating platform owners like engineers.

These data products should be available from a self-serve data platform service—think microservices, but for analytics data rather than applications. On these platforms, a data product is managed along its full lifecycle, while being enmeshed with other data products via relationships defined by code. This allows knowledge graph and lineage across the mesh, and provides the user with a frictionless experience as they look for data products.

The data should be discoverable, addressable, and trustworthy, so it needs to have an owner who is close to it, either as its source or as its main consumer. And so a data mesh is federated and domain-driven, meaning that ownership is the responsibility of the business domain it is associated with, rather than a central authoritative data management layer. The lines between the nodes are relationships formed by the users who access the data products from their respective domains.

Because the domains own their own data, accountability for data is also federated, with domain-specific experts responsible for its quality and integrity.

If the data mesh still seems a little abstract, consider a crude example. An analyst within a bank wants to understand correlations between, say, customer experience and economic events. The analyst wants to look at transaction datasets, social media datasets, and client records. In a centralized paradigm,



they would have to make a copy of all these datasets and download those.

Not only does the data quickly balloon to millions of rows and columns; the analyst may also struggle to trust entirely in its quality and integrity, as it has no clear owner and its timeliness may be suspect.

“In the big red blob, data tends to be stored for a long time. Data comes from various sources, which requires the blob to stay up to date. This is a tremendous amount of work. Why not just go to where the data is needed?” Carey says.

In a data mesh, the analyst can discover and access the datasets they need as products all in one place, on the self-serve platform. This is why one of the first things UBS did when building its mesh was to develop a browser, Carey says.

The analyst might make a copy to run a complex training algorithm on it, test out their idea and decide that their hypothesis was wrong, and they need to start again with different data.

“You wouldn’t want to [make a copy] against a production system, but with a mesh you could bring a copy over into a different node for a short period of time. And when done, the mesh can purge the data,” Carey says.

Also, the analyst can trust the quality and integrity of the data, because it is

overseen by the people who are most familiar with it. As the analyst pulls this data, they may create their own dataset, which then itself becomes a product and is available for later users. The relationships between the datasets—the ones the analyst has used and the ones they have produced—will be encoded as metadata. Metadata defines the “lines” between nodes, the relationships between one product and another. This metadata contains information about, for example, the size of a dataset or when it was last updated, but also crucially how it combines with other datasets.

This process of building the nodes and encoding the relationships between them isn’t a “free-for-all,” however, Carey says. “It takes work: Every line and every node needs to be well-defined as you build the mesh.”

For one, much of the data in an organization is sensitive. Perhaps it is confidential or falls under data protection regulation. Or perhaps a particular group of UBS customers would not care to have their data used in a particular way. Thus, every user in the mesh is permissioned to see only the data they have a business reason to see.

In UBS, the metadata about the nodes is stored in the firm’s DevOps platform, DevCloud. UBS launched

DevCloud, which is aimed at speeding up development cycles in the cloud, in partnership with GitLab in 2020. The node is about relationships, Carey says, and that meta description of the relationship between one node and another is what goes into DevCloud. When individuals have permission to access data, they can execute code to grab the data from the mesh and the code goes into DevCloud.

Carey says, however, that UBS does not put the data itself into DevCloud.

“We’ve made a large investment at UBS to create a unified development ecosystem because we treat everything as code, except for data,” he says. “It’s not best practice to put data into a DevOps repository. We don’t put data into DevCloud, but we put everything else there since everything is code.”

Carey says DevCloud is helping UBS implement a data mesh in a much stronger way.

“Without DevCloud, it is more difficult in situations where an analyst is trying to get data from three different places, three different nodes across three different businesses. When you have three different development platforms to get data, it’s not a great experience; for us, it’s all one experience,” he concludes. [WI](#)



Retooling repo

From brokers to trading platform providers, players in the repo market are focused on initiatives around reducing settlement fails, increasing automation, and streamlining operations. By **Nyela Graham**

When it comes to intelligent design and innovation, it can seem like the equities market gets all the love. Following that, not too closely, is the fixed-income market, which has been electronifying for years, but which still commonly relies on phone calls, paper, and interpersonal relationships to get trades done. The repo market, however, is an area that market participants say is long overdue for tech disruption.

A repo, or repurchase agreement, is a kind of short-term borrowing for dealers of government securities to use for raising short-term capital. A dealer will sell government securities to investors, typically overnight, and then buy them back at a slightly higher price the next day. For the seller and repurchaser of the security, such a deal is a repo; for the buyer and reseller, it is a reverse repo.

If that sounds simple in theory, it's not in practice. In the repo markets, frequent settlement fails, counterparty risk, and collateral management are persistent headaches. Settlement fails, for instance, when securities aren't where they need to be on the settlement date, may be reduced when various markets move to T+1 settlement, but are still common today. And counterparty risk is seen as a primary exposure in repo, as collateral could be made worthless by a default.

Nicola Danese, head of European fixed income at Tradeweb, says that despite efforts to modernize workflows and trading tools, most fixed-income traders seem to choose their old-school phones and messaging apps over disruptive products to make their trades. Repo, like other areas of fixed income, is a relationship business, and Danese says some traders still actively resist new interfaces, protocols, and platforms.

"The reason why people have always been concerned in the dealer-to-client space in electronifying repo is this is a relationship business," he says. "The concern was if I move into electronification, I'll lose the value of my relationship."

Still, resistance hasn't stopped tech vendors, venues, and even some dealers, from trying.

In the last year, projects and initiatives meant to digitize fixed-income trading by big names have stolen

"The reason why people have always been concerned in the dealer-to-client space in electronifying repo is this is a relationship business. The concern was if I move into electronification, I'll lose the value of my relationship." Nicola Danese, Tradeweb

headlines. Octaura, the product of a Citi-led bank consortium on loan trading, launched as a company in June, with its platform due later this year. The top three fixed-income venues by trade volume—MarketAxess, Tradeweb, and Bloomberg—confirmed plans to compete for the contract to deliver the regulator-backed consolidated tape for bonds in the EU. These developments followed what Coalition Greenwich, in a report last year, called an 18-month fixed-income "e-trading tailwind," spurred by office and trading floor closures due to Covid-19.

In an episode of the Waters Wavelength podcast earlier this year, Enrico Bruni, managing director of Europe and Asia at Tradeweb, called repo the bloodstream of bond markets.

"An efficient bond market is only as efficient as an efficient repo market," Bruni told *WatersTechnology*.

Bruni attributed the market's inefficiencies, complexities, and its slow trudge toward the 21st century to the "different flavors" of deals—bilateral versus tri-party trades, sell/buy or buy/sell repo, reverse or vanilla repos, and equity repos—seen in the space.


These nuances are what make one-size-fits-all solutions and one-stop shops few and far between. But from APIs to blockchain to smart order routing, firms such as Broadridge, JP Morgan, Tradeweb, and Bloomberg are throwing the technologies that transformed equities trading at the wall and hoping something sticks.

Automation, please

Tradeweb, for one, is investing in technologies that will automate workflows in the repo markets. The venue's most recent release is its Lifecycle Blotter, which supports lifecycle events for trading on-open, a type of trade involving repos without pre-fixed maturity dates. One example is where a deal can be terminated on any business day in the future by either party once they've given notice within an agreed period. The new functionality is meant to enable users to negotiate and trade actions on the platform such as partials, re-rates, re-prices, and close-outs.

"As central banks continue to dial back monetary support, we expect to see a return of more traditional liquidity sources for market participants. This should be a positive for the repo market and, ultimately, for the growth of the automated workflow that now exists to support it," says Tradeweb's Danese.

For its part, Bloomberg has centered on automation tools and APIs for its investments into fixed-income technology, and the data giant plans to launch an electronic trading offering, in collaboration with Euroclear and Sunthay, aimed at creating a front-to-back solution for bank-guaranteed repo. The forthcoming



service will have a dealer sit between clients to act as an insurance policy against either of the counterparties defaulting.

It's designed to enhance the bank's capacity for credit intermediation, and it allows end users and providers to connect via a bank intermediary, says Derek Kleinbauer, Bloomberg's global head of fixed income and equity e-trading.

Using proprietary standard master agreements from Sunthay, the offering will sit on Bloomberg's electronic trading infrastructure, which will support integration with each market participant's existing infrastructure. Euroclear will offer the post-trade services layer to bring tri-party services into existing custodial relationships.

"We've also seen an increase in the demand for automation, being able to execute either through a rules-based engine or through an API, where you can have that API plugged into your own internal systems and route via smart order router to multiple select venues," Kleinbauer says.

When all else fails, try blockchain

Last year, Broadridge launched its distributed-ledger platform for bilateral repo trades. The platform utilizes DAML smart contracts from Digital Asset to simplify the complex multi-


Horacio Barakat, Broadridge's head of digital innovation for capital markets, says the platform digitizes the repo process from front to back using smart contract technology, which allows counterparties to sync with each other, starting at execution. It creates a shared workflow for counterparties, which can then be linked to settlement through the tokenization and digital presentation of collateral.

The platform is interoperable and leverages existing market infrastructure, Barakat says. The interoperability is important, he adds, because it should increase adoption, particularly by encouraging risk-averse institutions to add something new into their workflows, rather than need to overhaul them. Down the line, there will be the possibility of interoperability between other ledgers and platforms as more come to market.

Broadridge is not alone in utilizing blockchain technology for the repo market. In 2020, JP Morgan launched

"Because it's all happening on the same shared ledger, we can leverage things like smart contracts, and we can get much more precise about the length of time that our clients actually need to borrow these funds for." Tyrone Lobban, Onyx Digital Assets

party workflows in the repo market. It also uses the VMware Enterprise blockchain platform to provide the underlying cryptographically secure distributed-ledger network. Last month, Societe Generale went live on the platform, joining UBS, which onboarded in August 2021.



the Onyx Digital Asset blockchain-based network and its intraday repo solution, which came to fruition following years of research and development at the bank around blockchain use-cases.

“We actually began our blockchain work at JP Morgan in 2015 and when we began, quite like many other financial institutions and companies at the time, we were looking at where blockchain makes the most sense to apply,” says Tyrone Lobban, head of blockchain launch and Onyx Digital Assets.

The bank was looking to identify use-cases in which the then-fledgling technology could create new products or bring efficiency to existing ones. During those years of research, Lobban says the bank executed close to 100 proofs-of-concept, some of which included matching trades between brokers and custodians and corporate actions processing.

In 2017, attention shifted to the tokenization of bonds by representing them on a blockchain. Through

a partnership with the National Bank of Canada in 2018, JP Morgan issued a \$150 million, one-year floating-rate Yankee certificate of deposit using blockchain technology.

The Onyx platform would be born in October 2020, building on the idea that a traditional asset could be represented on-chain.

“The idea here is that we should be able to represent any type of traditional asset on-chain and start to use a blockchain for faster settlement, more transparency around that settlement, reducing fails processes, and having stronger guarantees around the delivery of assets in exchange for payments,” Lobban says.

The intraday repo solution launched two months after Onyx, allows JP Morgan clients to borrow funds intraday from the bank on a secured basis, while putting up collateral in the form of tokenized US treasuries.

“Because it’s all happening on the same shared ledger, we can leverage

things like smart contracts, and we can get much more precise about the length of time that our clients actually need to borrow these funds for,” Lobban says. Instead of entering a repo transaction that has a standard overnight maturity, a client could theoretically borrow \$1 billion for three hours and return it before lunch.

In May, the bank launched the Tokenized Collateral Network on Onyx Digital Assets. Lobban says the work around the intraday repo solution led to a broader thesis around generally improving collateral and representing different types of assets on a blockchain.

“In the collateral markets, there’s not just the challenge around the settlement timeframes alongside the need to move pieces of collateral, which you are looking to deploy in lots of different places, whether through tri-party agents or through different bilateral counterparties,” says Thomas Pikett, vice president of trading services product management at JP Morgan. “It’s also about being able to make the maximum use of your asset pool available as collateral.”

Like many firms and providers that embrace emerging technology, JP Morgan is aiming to create a service that clients can easily acclimate to. A representation of an asset should be as close to an asset as possible, with the accordant reference data, such as an Isin, Pikett says. The record of ownership should be obvious and as easy to understand as a traditional asset.

“It’s quite powerful in itself, because you can get the benefits of blockchain, you can get the benefits of tokenized assets, but you’re not needing to rebuild systems to understand a new funky asset,” he says.

As the more bespoke components of fixed-income markets have proved resistant to change, sources agree that a full overhaul of the existing workflows in the market is not welcome, nor possible. But if it’s true that an efficient bond market is only as efficient as its repo market, then change is necessary and coming. Keeping familiar processes and frameworks in place may coax even the most reluctant of traders into the digital age. [wt](#)



From pharma to finance: Cracking the DNA of data management

Enterprise data management has traditionally addressed any aspect of financial data across an organization. But as investment firms' portfolios of enterprise data broadens in definition to include other types of non-financial data, EDM projects must also expand to embrace new data types. By **Max Bowie**

As investment firms seek an edge to help them make better and more profitable decisions, they're delving deeper into the data behind industries, companies, and their products to perform more in-depth analysis. However, in many industries, while detailed data on a company and its business areas may be available, it's often not available in a format that is easy to integrate with price data, financial data, or reference data.

So for investment firms focused on a specific industry or theme—such as biotech or ESG, for example—there aren't yet established identifiers or ontologies that allow them to combine the data that they believe will provide extra insights with the functional financial data residing in their enterprise data management systems. And while the efforts of industry asso-

ciations and the vendor community are giving hope for future data interoperability, this leaves those wanting to take advantage of that data today either plowing time and money into custom data mapping and integration projects, or trying to squeeze the “square pegs” of new data types into the “round holes” of existing technology platforms.

Boston-based fund manager RA Capital invests in public and private companies across the medical, healthcare, and life sciences fields. The fund looks at specific factors that could impact healthcare, pharmaceutical, or biotech companies, such as drug names, formulations, and clinical trials. The problem is, that data doesn't easily translate into numerical values, nor does it come with the kind of identifiers or descriptive reference data

that would allow a firm to easily map it to a company name or stock price.

The firm was founded in 2001, and until now had let its research team manage the process of sourcing and correlating data to drive investment decisions manually. As the firm grew its assets under management (today it manages over \$10 billion), the volume of data it subscribed to increased. This included news feeds, as well as niche content sets that track relevant events such as drug trials and approvals, and attributes of individual drugs, in addition to company-related data such as information on members of a company's board of directors.

As a result, the need to invest in a proper enterprise data management system became more transparent to executives at RA Capital. Not all of these datafeeds easily lend themselves

to traditional market data platforms, and would commonly be managed in an ad-hoc manner in spreadsheets or an analyst's own database, says Chris Caliri, chief information officer at RA Capital.

"Our data was becoming unmanaged and uncontrolled because we didn't have a central repository. We noticed things like company records going stale, or missing data. For example, a corporate action could cause issues and we would have to ask why certain values were wrong. If we were a \$1 billion shop, we most likely wouldn't need an EDM platform; we'd still be doing all this in spreadsheets and siloed databases," Caliri says.

But even though the firm was spending thousands of dollars on datafeeds, it wasn't seeing the full benefit because it couldn't efficiently access that non-financial data and match it to relevant market data.

So, in 2020, the firm began evaluating EDM platform providers, including UK-based data technology vendor Xenomorph. The firm kicked the tires on about a dozen providers before picking Xenomorph. Each could do some of what RA Capital needed, but would either be unable to handle specifics within the data and the mappings between specific drugs and their manufacturers, or be unable to get data in and out of their platforms quickly. Or, it would instead turn into a longer-term consulting project. Caliri says Xenomorph could demonstrate on the fly how to recreate the firm's workflows in its data model.

"We wanted to implement a platform like Xenomorph that could manage all data and feed it to users and/or downstream applications. We needed to get away from our manual processes," he says.

After running a proof-of-concept in late 2020, Xenomorph began officially working with RA Capital in early 2021, and by September rolled out the first phase of its platform. That first project involved understanding the firm's needs and the data modeling required, connecting to the firm's data suppliers, matching data between external vendors and internal systems, and helping it generate analysis and reports based on the data, says Naj Alavi, New York-based president of Xenomorph.

Part of the challenge was mapping the key interrelationships between different categories of data defined by RA Capital that all relate to the same drug or manufacturer—not just a company name or drug brand, but separate events, indicators and milestones that relate to all aspects of a drug and its development, says Mark Woodgate, who founded Xenomorph in 1995.

"The complexity of physically storing the data is one thing, but with all of those complex interrelationships, a standard security master just wouldn't fly. You would need to create a whole new schema," Woodgate says. Because

"Our data was becoming unmanaged and uncontrolled because we didn't have a central repository. We noticed things like company records going stale, or missing data." Chris Caliri, RA Capital

Xenomorph's platform is data-agnostic and able to support any type of data, Woodgate says it was able to match and merge data from multiple external vendors each supplying different data elements with the firm's internal systems to link data to companies and drugs, even for companies that aren't publicly listed.

"The security master part was fairly straightforward, especially on the publicly traded equities side," says Ernesto Gonzalez, senior business analysis manager in RA Capital's IT department. "It gets more challenging to match and merge on the private companies side because there are no identifiers. So, we partnered with Xenomorph to build out rules capable of merging data from different vendors and various types of entities—for example, companies and drugs from two different data providers."

Xenomorph's EDM+ platform has been in user acceptance testing at RA Capital during phase one of the implementation. At that point, having onboarded all of the firm's data, maintaining it becomes an incremental process of adding new companies and

drugs as they become of interest to the firm, he says.

"We're now starting to leverage what we've built. Even though we're not technically live, we've had a number of ad-hoc projects where we were able to provide consolidated data from a number of other vendors to the research team, whereas previously they would have been tasked with manually maintaining the spreadsheets that we're now doing systematically," Gonzalez says.

"Now we have a central and systematic workflow of merging and matching data—so a lot of that manual work goes away," he says. "So, whereas the research team has previously been focused on data gathering, now they can focus more on analysis and validating the results."

Something else that goes away is the potential for error that exists within manual processes. Caliri says replacing spreadsheets with a proper platform for data management not only reduces potential errors from copying and pasting data in Excel spreadsheets, but also eliminates the challenge of extracting data from a spreadsheet once it's been entered into one.

Even then, there was still a lot of manual effort involved in setting up the platform.

"Because of the sheer volume of data we ingest, it was a big effort," Gonzalez says. "Xenomorph automates a lot of the matching and merging, but there are still a lot of manual reviews that need to happen."

Defying definition

RA Capital isn't alone. Other firms attempting to use granular and non-traditional data to fuel investment models are encountering similar challenges. Xenomorph's data-agnostic approach allows the vendor to easily onboard new data types—for example, the vendor has already run a proof-of-concept using wind turbine data for clients in energy markets. Yet, despite the complexities involved, some prefer to tackle these challenges in-house.

A good example is ESG data—a rich tapestry of information spanning myriad topics and data points, much of which defies the tabular and numerical formats of traditional financial data. Not only is



the relevant information more likely to be textual and require translation into any numerical value, rating, or score, but each firm may place a different emphasis on each.

RadiantESG, a San Francisco Bay Area-based investment fund, focuses on using ESG factors and models to identify companies that represent good investments and that make a meaningful effort to incorporate ESG into their day-to-day business.

“We try to combine fundamental and ESG factors,” says Mauricio Bustos, head of data and technology at RadiantESG. “So, we’re looking for companies with a solid financial foundation, but on top of that, companies that demonstrate a credible intent to push forward in the three areas of ESG.”

However, while there’s plenty of ESG data available in the market, obtaining data from its original source in any kind of standard format can prove challenging for firms that see value in performing collection and analysis in-house on data including financial filings, lawsuits affecting a company, corporate and social responsibility reports, data from aggregators, and information from NGOs about controversial industries or companies.

Bustos says regulators do not require companies to produce this data in a standardized format, which leads to firms needing to rely on third-party vendors, NGOs and public sources to do that. Firms also need them to extract information from, for example, 10-K and 10-Q filings about what a company is doing and to detect a company’s credible intent to contribute to these three pillars. Yet another issue is that ESG data doesn’t adhere to traditional delivery formats in the same way as other market data such as prices or corporate actions.

“It doesn’t come in a firehose, but rather trickles in,” Bustos says. “So, as it does, you have to build up a full picture of a company, like building a mosaic.”

That said, while corporate and social responsibility reports aren’t *required* in equities markets, plenty of companies already report that information voluntarily, so Bustos says it has been fairly easy to map existing data to its security master. RadiantESG made the strategic decision to run its data management in-house.

“Reference data management is a very important element that we need to own completely in order to have confidence that our model is working the way we expect it to,” Bustos says. “To make sure

we’re talking about the same company, reference data integration—and making sure that we’re integrating that data in a consistent way—is a critical component.”

Normally, a firm would rely on an off-the-shelf EDM platform to piece together that mosaic of data and map it to market and reference data about a company or instrument. However, RadiantESG built its own, combining cloud technologies for storage with natural-language processing tools to aid with details of text extraction, such as for entity identification, sentiment detection, and keyword evaluation.

Bustos notes that the group that built RadiantESG came from quant shop Rosenberg Equities, which was part of Axa Investment Management. At that firm, he says they built everything from scratch.

“The cloud has been a huge benefit in terms of managing these large datasets. Because there are no standard formats, data needs to be extracted from free-form text—and that text takes up a lot of space, because anything we want to extract data from needs to be stored somewhere,” he says. “So, being able to rely on the cloud to store huge amounts of data has been very beneficial and affordable.”

RadiantESG is the exception to the rule in building its own platform to manage its data assets, but not in terms of using custom tools for handling custom datasets. The complexity of bending data models to accommodate unusual data types is something that most firms and vendors have trouble with, says Ethan Shen, CEO of Crizit, whose Periscope platform monitors enterprise data licenses and usage.

“Reference data management is a very important element that we need to own completely in order to have confidence that our model is working the way we expect it to.” Mauricio Bustos, RadiantESG

“Where we’ve seen non-traditional data sources, they are typically consumed by bespoke internal applications. Most finance deployments of EDM platforms are built around a data model that represents companies, tradable instruments, and events. Once you stray outside of the standard models, customizing large EDM platforms typically is much more painful than just building a custom application,” Shen says. “For the EDM platforms typically used in finance, the models are quite rigid. Customizing is possible, but often prohibitively slow and difficult.”

But as the use of non-traditional datasets widens, managing them in EDM systems—and monitoring usage and cost—will become essential. TRG Screen, for example, already supports the ability to define different data types in its Optimize Spend data inventory and spend management platform.

Richard Mundell, chief product officer at TRG Screen, says the vendor has customers tracking market data-adjacent categories such as research by brokers and other third parties, as well as more niche publications. All these information sources tend to fall to the market data teams to manage. Through Optimize Spend, clients can, for example, create their own user-defined fields and use them to track individual data elements that comprise ESG information.

“Bear in mind that they’re not storing the actual data in Optimize Spend, but instead they’re tracking the necessary information to track the subscription—so that’s all the details of the vendor, the contractual terms, the services covered under that contract, how those services are consumed by the firm (including platform and delivery mechanism, if applicable), and onward from there, the inventory of who has what and the cost allocations,” Mundell says.

TRG Screen needed to build extra features into Optimize Spend to account for the nuances of specific non-traditional datasets. For example, Mundell says, firms using expert networks such as Gerson Lehrman or AlphaSights tend to pre-pay for a set number of meetings and consume that balance over the duration of the agreement. Clients can define their pre-paid amounts, and the system will track and report their remaining balances. Mundell says the company will add similar capabilities as clients expand their use of the service to new expense categories.

Standards to the rescue

But the fact remains that without standard definitions and identifiers for data types that don’t yet have a formal structure, managing them will be up to bespoke in-house implementations and costly consulting projects that may solve one firm’s issues but can’t be scaled industry-wide to gain greater efficiencies.

However, standards could soon be on their way, thanks to efforts by the EDM Council, a reference data industry association, which—as part of its effort to expand its reach outside of the financial arena—is developing semantic ontologies for other industries that could create standards to also benefit investors monitoring those industries and companies within them.

Around four months ago, the EDMC began a partnership with major pharmaceutical industry players, including Bayer, Merck, Roche, GSK, Johnson & Johnson, and others, to develop a pharma-specific ontology, dubbed Identification of Medical Products (IDMP), which identifies everything from individual substances to units of measurement and maps them together. The primary driver for this

effort is to make it easier for consumers to identify the same drug in different countries and jurisdictions, which may have different names or different formulations. However, it could also make it easier for companies like RA Capital to identify and compare the offerings of different biotech and pharmaceutical manufacturers and provide a more structured basis for apples-to-apples financial analysis where that hasn’t previously been possible, or has only been possible by performing custom projects.

“It’s a mess in the pharma industry, and that mess is also what we’re being approached about for other industries, like ESG and agriculture. The EDM Council is getting involved in creating ontologies for these as well,” says Mike Meriton, co-founder and COO of the EDMC. “For example, we’ve formed a 100-company working group on ESG data, looking at how companies report, how vendors get that data and publish it, how consuming investment companies use it to make decisions, and how standards bodies standardize it.”

As a result, EDMC is now being asked to apply its expertise to create a climate ontology, Meriton says, so that information traditionally stored in different databases and referred to using different terminologies in each can now be harmonized so that a standard identifier can be mapped to the same data in each usage, allowing firms to aggregate separately stored data without having to consolidate it into a single, new database.

“So, whether you’re an investment analyst doing drug discovery or a pharmacist helping someone find the right drug, an ontology ... allows you to identify, correlate, and match data at scale, whereas today that uses multiple different standards,” Meriton says.

Once standards are structured in a way that they can be applied to any company, product, or industry, then data on all those elements will have the same attributes as other, structured types of financial, market, and reference data, and firms will be able to more easily manage any potential data input alongside—and seamlessly integrated with—the data they use to make trading decisions, without the need for major, custom integration projects. [wt](#)



Cloud and APIs begin to (slowly) permeate the post-trade space

As financial firms turn their attention toward modernizing the back office, how they approach these projects comes under new scrutiny. By **Wei-Shen Wong**

Front-office trading platforms are increasingly getting cloud makeovers. The same is true for risk systems sitting in the middle office. The next frontier for cloud evolution would logically be post-trade and the back office. But the back office is an area that is notoriously underfunded and contains a hodgepodge of legacy systems that connect to myriad different applications—whether that’s data coming into the firm, or data moving from the front and middle offices into the back office, and vice versa.

However, as more banks and asset managers embrace the cloud and want tools delivered as a service or as a managed service, more vendors are responding to demand in the post-trade space to help customers figure out where they should focus their attention as they migrate systems to the cloud.

Firms could start, for example, with position management, says Danny Green, head of international post-trade at Broadridge, as this is an area where multiple systems are performing the same function. “In theory, you could deliver a general position management component, and then ultimately switch it off from all of the multiple systems,” he says. “By doing that, you get to have a global position management capability quite quickly. So instead of replacing the entire system, you start with replacing position management, but everywhere.”

That way, firms can have visibility of global positions in real time, and provide that data via APIs from the back office to the front office. “That would be a function you could deliver as a single component, for example. That then shrinks the overall foot-

print of your legacy systems and gets you on that journey of slowly replacing them,” Green says.

James Marsden, managing director and head of post-trade business for Asia-Pacific at Broadridge, says it’s easier to put something new in the cloud rather than to migrate a legacy platform to the cloud.

“Sometimes in a legacy platform, the code has grown over 20 or 30 years, and nobody quite knows what it does. But when you implement a new platform, or a new component—like position management, for example—we can implement that as a cloud-enabled component. And it doesn’t have to be everything in one shot.” James Marsden, Broadridge

“Sometimes in a legacy platform, the code has grown over 20 or 30 years, and nobody quite knows what it does. But when you implement a new platform, or a new component—like position management, for example—we can implement that as a cloud-enabled component. And it doesn’t have to be everything in one shot,” he says. He recommends finding areas that logically lend themselves to tools that don’t have so much coding drama underneath.

Green says Broadridge has been looking at transaction capture, a middle-office function, since its acquisition of Itiviti. If firms put a layer between the front office and all their post-trade systems by looking first at the middle office, they could improve how they send trade confirmations to

clients, and how they communicate generally with clients, he says.

“We’ve been doing a lot of studies of that kind of front-to-back relationship, and that whole transaction-capture, middle-office piece is another good place to start,” he says.

The key, Green says, is to plan carefully before deciding on any innovation program. “What does the future look like for you? That could include things like data, the use of APIs and other technologies such as cloud and distributed-ledger technology (DLT). Once you’ve got the target operating model of what you want to achieve, now let’s work with you to create a roadmap of how you can achieve that,” he says.

The API play

Many banks’ core systems reside in the back office, making ripping out and replacing those systems a nightmare.

Gurvinder Singh, CEO at New York-based trading, risk, reporting, and data management solutions provider Indus Valley Partners, says that, where possible, firms should look to lift-and-shift monolithic applications from on-premises to the cloud. From there, the key is to re-architect applications and services to a microservices model so that they can more easily be transitioned, should the need arise, in the future.

“The next stage [after the lift-and-shift] is to start carving services out that make sense, while preserving legacy monolith cores where there are no functional or performance challenges,” he says. “This is the design pattern that has been successfully implemented by many enterprise B2B firms, and is the way to go for all.”

These systems have been built over many decades, with some still running on the Cobol programming language.

“They’re older, they’ve got a lot of logic and nuances that are not as easily migratable to cloud infrastructures,” says Neelesh Prabhu, managing director of architecture and enterprise services in information technology at the Depository Trust & Clearing Corp. (DTCC).

This means that firms approach these modernization efforts incrementally, so there won’t be a big-bang approach to innovation in the space. Prabhu says, though, that it’s relatively straightforward to move systems that sit on the edge of the overall IT ecosystem to the cloud.

“That is the web front ends, and the systems of engagements that banks and larger financial institutions have built. Those are relatively modern as they’ve been built in the last 10 to 15 years, and can be easily adapted to the cloud,” Prabhu says.

This is where APIs come in—provided they’re done right. With APIs, banks are moving small processes/pieces to the cloud, and in doing so, de-risking the programs in question.

“On the front end, there is heavy migration. On data, there is a lot of interest and migration. And on the core systems side, there is migration, but firms are choosing to do it in a way that’s mindful of the risk that moving some of these systems to the cloud may bring to them,” Prabhu says.

He says APIs provide the ability for firms to connect the logic and the functionality provided by a particular system with other systems, but at the same time, lets the teams hide the internal details of how those functionalities are implemented.

Say there’s a system built on mainframe technology. The first thing to do is build an API layer around that system to create an endpoint for all the other systems connecting to it. Once that’s done, the firm can replace the internal technology with a more modern system.

“It’s the idea of using the API as a construct of encapsulation and the cloud to bring functionality quickly,” he says. But the key to these platform conversions is for both the business and technology teams to work together. The design of the API needs to make sense in the business context.



“As these systems talk to each other—it’s not just one connectivity point, but that the structure is built with business understanding that they can stand the test of time,” Prabhu says. “The structure you’re putting in place is foundational to future connectivity points which could exist or come into play.”

Born free


Innovation in the capital markets is often cyclical. Pre-2008, banks were all about building proprietary systems; after the financial crisis, because of a flood of new regulatory demands, banks cut IT staff to the bone and leaned heavily on third-party providers to lower costs and improve margins.

The idea of leaning on a major public cloud provider like Amazon Web Services (AWS), Google Cloud, Microsoft Azure or IBM Cloud for key trading and data management needs was dismissed on sight. But slowly, banks—and even asset managers—have begun to

consider cloud. New regulations led to firms having to suck in and store increasingly larger amounts of data. At the same time, new datasets became available as, vitally, the ability to store and run compute on massive datasets became more viable as public cloud providers made these services more cost-effective and improved time to market on new tools.

This gave birth to the fields of alternative data and software-as-a-service models. More data and the availability of tools to analyze data—and, thus, find previously unforeseen correlations or areas of risk or productivity enhancement—led to firms wanting not only more data, but they wanted context wrapped around that data.

This is a major reason why firms want to modernize their post-trade processes: It provides more and better-structured data from which to derive insights, rather than leaning on creaky legacy platforms. It provides the potential to deliver alpha and reduce risk.



As more vendors look to shift their legacy platforms to the cloud, several startups have come to market looking to jump ahead of stalwart vendors because their tools are born in the cloud and utilize new data delivery systems, most notably, APIs.

“Cloud creates the infrastructure for allowing the automation of expensive and risky workflows, creating a situation where staff can focus on areas to create value,” says Brad Bailey, head of market intelligence at broker-dealer Clear Street. “APIs are another key tool to facilitate the exchange of information.”

Clear Street launched in 2018 and its mission is “to build better infrastructure to improve market access for all participants.” As a startup, it has its sights set high and there’s no way of knowing whether it will succeed, but if the company can figure out a way to streamline the post-trade space, that could give it a leg up on competitors

that are in the process of migrating legacy systems to the cloud.

“Solving the fundamental problems in the post-trade ecosystem requires a native rebuild for core problems to be addressed,” Bailey says.

Then there is RQD, a company that obtained a limited clearing license in 2018. Instead of establishing presences at Equinix datacenters, it decided to go all-in on the cloud. RQD uses Microsoft Azure and other Microsoft technologies, such as SQL Server and .Net.

RQD COO Nicolas Louis says the company had the luxury of starting the process from scratch.

“When you are a clearing firm that’s been around for a long time ... you may have thousands of processes running against your [on-premises] systems every day—whether it’s a margin calculation, updating a position, or building a report for Finra, those systems are working non-stop. So now, people have been saying, ‘You need to main-

tain these systems, and that’s expensive and hardware can fail. Let’s move to the cloud,’” Louis says.

But this is where firms encounter a challenge because moving features to the cloud is easier said than done. One reason for that is they are moving it bit by bit, taking very specific pieces of their process, isolating it, and moving it to the cloud while everything else runs on-premises.

“So you have this hybrid setup, which sounds like a great idea, but all you did was add more points of failure,” Louis says. “For example, you may lose the connection to your cloud, and end up with a worse position, with processes running on different environments but not able to communicate with each other.”

Stability, resiliency, and the availability of post-trade systems are hugely important. If the lights go out, it could result in duplication of transactions or settlement failures for which firms could be penalized. So they end up running both systems in parallel.

“You have the new environment running in the cloud for a specific need, and you have a shadow back-up environment running on-premises that still needs to be supported,” he says.

The arrival of startups like these, paired with the modernization efforts at established players like Broadridge, DTCC and Indus Valley Partners,

“Solving the fundamental problems in the post-trade ecosystem requires a native rebuild for core problems to be addressed.”
Brad Bailey, Clear Street

shows that there is a shift underway when it comes to cloud and post-trade. As ever, the back office will lag behind the front office, but nevertheless, change is coming. Firms will have to consider what they aim to achieve in their modernization process, and then only look at applying or using technologies such as cloud and APIs to help them get there. Applying the hammer-looking-for-a-nail ideology will not fly. **WT**



After switch to Google, Symphony begins data migration for clients

Symphony decided to first migrate client data to Google from Amazon Web Services, and then focus on applications.

By **Melissa Fleur Afshar**

In financial services, the move to the cloud is inexorable for banks, asset managers, exchanges and vendors alike. The challenge is in figuring out the right path to take.

Trading workflow and communications platform provider Symphony Communications Services initially signed a deal with Amazon Web Services (AWS) as its primary cloud infrastructure provider. But after closer examination, executives deemed AWS to be too expensive and not the best partner for its desire to deploy AI-driven tools.

So in July 2021, Symphony announced a partnership with Google. It intended to begin migrating to Google Cloud Platform (GCP) in Q4 2021. “The process is expected to take about a year to complete, and will see Symphony’s core platform run in Google Cloud, with some associated services that may still need to run on AWS,” *WatersTechnology* reported at the time.

More than a year in, the migration is “on point, on target and on time,” Brad Levy, CEO of Symphony, tells *WatersTechnology*. But changes have had to be made.

“When you get into the reality versus the plan, if you didn’t change anything, then you’re probably not right. ... As we’ve gone into it for real and spoken with their engineers in depth, we’ve learned a few things and changed our plans in some ways,” Levy says. “That shouldn’t impact the timing, but will impact ultimately how and what we do, and when.”

This month, Symphony began migrating first smaller and then midsize users’ data into production on Google, says Dietmar Fauser, Symphony’s chief information officer.

When the partnership was initially signed, the idea was to migrate applications and data in parallel. After getting a handle on the project that was to come, and getting closer with Google engineers, Symphony decided to first get the data right on Google’s cloud, while

“When you get into the reality vs. the plan, if you didn’t change anything, then you’re probably not right. ... As we’ve gone into it for real and spoken with their engineers in depth, we’ve learned a few things and changed our plans in some ways.”

Brad Levy, Symphony

also moving over some easier microservices, Fauser says, though the goal was never to “lift and shift from AWS to Google,” but rather to first establish, where possible, which processes and workflows could be automated.

“Currently, we are ending the full-blown testing of data migration from AWS to Google. That was a bit of a

change in the project—to go with data migration first—because we have some technical debt with some old databases running on AWS. So we consolidated Bigtable and MongoDB running on Google. We are currently finalizing this. [This month], we will start the data migration in production onto Google,” Fauser says. “So overall, pretty good progress.”

Previously, Symphony was running multiple instances of the Apache HBase open-source database on AWS. But to improve performance and to save on cost by reducing the amount of human resources needed to oversee all those instances, the vendor decided to switch to a large, shared instance on Google’s Bigtable NoSQL database. He notes that Bigtable also provides a better geographical footprint and redundancy.

“We changed the approach a little bit in the sense that we started with the data migration first, which was not the initial plan, but we were running into an end-of-lifecycle issue with HBase. So we changed slightly and did this first, which had impacts on the initial plans. It’s a relatively complex





“

“Currently, we are ending the full-blown testing of data migration from AWS to Google. That was a bit of a change in the project—to go with data migration first—because we have some technical debt with some old databases running on AWS. So we consolidated Bigtable and MongoDB running on Google. We are currently finalizing this.” **Dietmar Fauser, Symphony**

thing that requires more coordination with customers than what we initially thought; it’s a little bit slower than what we thought at the beginning, but overall it’s progressing fine,” Fauser says.

Symphony’s video meeting dataflow service has been migrated to Google, as has its base infrastructure. Now Symphony is moving client data over to GCP from AWS.

Levy says the migration process will last into next year for the company’s larger tier-1 accounts. While the project isn’t necessarily onerous from an engineering standpoint, the hand-holding and changes to services require time and care.

“None of this is turnkey—the complexity sometimes comes in the configuration and how [clients] deploy and manage keys, etc. So we have to work with them very closely to do the migration,” he says.

Cloudy skies

Major cloud migration projects are not new to Levy. Prior to joining Symphony in 2020—first as president and chief commercial officer, and then as CEO in June 2021, replacing David Gurle—he was CEO of MarkitSrv, a platform that has been described as the nerve center of the post-trade market in derivatives. While there, he oversaw one of the

more ambitious projects in the capital markets, as the IHS Markit-owned company launched TradeSrv, a cloud-based platform that would eventually house its other platforms, including MarkitWire, Markit Trade Manager (MTM) and DSMatch.

As with any ambitious project, there were bumps along the way and lessons learned.

“I’ve been involved in a lot of migrations like this where it’s an applications and data,” Levy says. “You often get very far down the road and then you realize something is wrong with the data. So, we now work with the data first and then more of the applications second. Normally you’d get the application ready and then put the data in it, but this will make it a cleaner, more confident migration at the end because we know we’ll have the data right by the time we start making sure the applications work.”

Levy and Fauser hope the benefits will be tangible relatively quickly.

They say that by switching to Google from AWS, in addition to improving performance and stability while cutting human labor costs, Symphony will be able to go to market with new products more quickly. Additionally, Fauser believes Google has a “different and much easier way to use the global network, which simplifies our operations and connectivity a lot. Data storage, and data management in general, are more advanced on Google, in my opinion. It allows for cost-efficient management of very large datasets. We use these capabilities for the databases but also for the [business intelligence] data that we collect and make available to customers on a large scale.”

While moving client data to Google is the goal right now, down the line,

“None of this is turnkey—the complexity sometimes comes in the configuration and how [clients] deploy and manage keys, etc. So we have to work with them very closely to do the migration.”
Brad Levy, Symphony

the partnership with Google will help it to expand in AI capabilities, specifically around machine learning and natural language processing. This will tie in nicely, they say, with the opening of a new R&D center based in France.

“As Google’s suite of apps is cloud-native, collaboration-enabled and used by many schools and corporates, we expect more demand over time to

weave into workflows and connect via APIs/extensibility. Other capabilities in data science, AI, engineering and identity are also being considered,” Levy says. “AI is something we are trying to be thoughtful about and produce a platform that can deliver AI outcomes for our clients. We have an R&D effort we just launched in France that is very data-centric and ultimately looking to AI. So I’d say we’ve achieved the building blocks for this in the last year.”

But when it comes to using Google to build new analytics, encryption and AI tools, “it’s still early days,” Levy says.

Symphony is currently targeting mid-2023 for the full migration project to be completed—but as we’ve seen with the cloud, plans can change. **wt**

Google–Symphony: Something to see here?

Anthony Malakian wonders if there are any tea leaves to be read as a result of Symphony migrating its platform from AWS to Google Cloud.

Editor’s note: This column was originally published in August 2021. We reprint it now to provide some context as to the significance of this pairing.

In July 2021, it was announced that Symphony was switching from Amazon Web Services (AWS) to Google Cloud. As executives from Symphony told Max Bowie, there were a few things that led to the signing. For one, Google was a cheaper option, they said. But beyond price, they believe Google will help the collaboration platform provider build out its data-sharing capabilities through the use of the tech giant’s AI tools and big data analytics services.

But what the execs made clear was that this is not purely an infrastructure play.

“We see infrastructure as a commodity. So the focus of this is what will come from working with Google beyond infrastructure,” said Dietmar Fauser, chief information officer at Symphony. “What excites us is that in the discussions with Google, it became clear that there was the potential for more, beyond just classical cloud provision—for example, around data sharing, where Google is one of the historical leaders in the space.”

And Brad Levy, Symphony’s CEO, had this to say: “We’re focused on new issue markets because they are more of a ‘greenfield’ scenario—there’s a lot of room for digitization and because they involve a lot of confidential data. Transactions data needs to be widely shared—with the parties involved in a deal, securities services providers, fund administrators, and custodians. Right now, a lot of that information is held in deal rooms and in emails. Our goal is to make sure that federated firm-to-firm transaction can happen.”

In Q4, Symphony will begin the process of migrating its platform to Google Cloud, with the project expected to be completed about a year thereafter. Fauser said that beyond cloud infrastructure needs, the company is seeing that more of its users are looking to lessen their over-dependence on Microsoft Office and are increasingly incorporating Google Sheets, Docs, and Gmail into their investment workflows.

I’m skeptical of that claim. The idea of moving away from Excel spreadsheets is a conversation I’ve heard about for at least a decade, and I’m sure it didn’t start there. But fair enough. The point is that in the primary markets—where there’s always been a great reliance on spreadsheets and emails—Symphony believes that Google can help streamline communication workflows.

If I were a betting man—and I am—I’d put a sizable percentage of my stake on Google, AWS, or Microsoft as being the company that one day gobbles up Symphony, as one (or probably all) of them looks to make a more concerted capital markets play.

I recently asked Gurle about just that. Here’s what he had to say: “Zoom didn’t exist really for our world, and now it’s a verb. Slack gets bought by Salesforce. Teams is rolling out [new services], and Microsoft owns the stack pretty much. Google is trying to up their game in cloud and become more relevant in the workflows. Amazon has the big presence but doesn’t do much more than that in our space. . . . What is the right thing for the long-term industry in

working with these very large technology companies that could be incredibly helpful, or so big that they may not know enough about you to care in the right way? We’re a bit of that bridge in bringing them in and managing them in a productive way. They will be disruptive in some spots and helpful in other spots, and so we’ll wait.”

I asked a long-time fintech CEO about the idea of a Big Tech company buying Symphony. It should be noted that the exec is not a fan of Symphony, but I do trust their opinion and I figure it can’t hurt to get an outside perspective in this column, as this person’s company does not compete with Symphony.

They say that they do not believe that a financial services company will buy Symphony, but they do not believe it will be a giant that buys them, either. They view this latest move as the natural fit simply because Google is already an investor in Symphony (making its initial investment in 2015), and argue that Google needs to start adding some big scalps in the capital markets arena to keep up with the other major tech companies growing in the space.

“Amazon is the default player in financial services, and Microsoft is coming up very fast because everyone is very comfortable working with [Microsoft Office] 365, and Azure has done some good things around security,” they say. “So this is Google’s way of forcing some financial services adoption of its cloud. That’s what happened here: We invested in you, you better get on our cloud. I don’t think this is a precursor to Google acquiring them, but I could be wrong.”


Fair enough. So if this announcement isn’t a precursor to a deal (and I’m not sold on that, for what it’s worth), then here are the questions that I have, which I suspect some of you may also have.

At the top of the list is this: Why didn’t this happen sooner—you have one of the major cloud infrastructure providers as an investor; why not use them? And why wasn’t Google pushing it? We didn’t get much of an answer on that front. Maybe the prior deal with AWS was a leftover from the original setup of Symphony or even from Perzo, which was the precursor to Symphony. Or, maybe the original backers were directly familiar with AWS, as Google hadn’t made a big capital markets push prior to 2015.

There’s no doubt the pair is strengthening their ties to each other. The cloud deal is strategic, and the additional capabilities they can get from working more closely sets them up for doing more in the future. It also allows Google to test the waters in a more active way, letting guinea pigs try on some high-tech wheels for size, and the deal gives it a much broader exposure within financial services. So this could be just an opportunistic quid pro quo, but when you have millions of dollars sunk into a platform, it makes sense for Google to put some effort into protecting (and hopefully growing) that investment.

I think what happens next will be a strong signal. First, who else does Symphony buy/ally with, and for what purpose? (Symphony announced it had acquired StreetLinx, but that’s not the kind of needle-moving deal that I’m talking about.) Second, will Symphony raise more money, and if so, from whom? Third, will any existing investors try to cash out, or will others buy out their stakes—essentially, will it turn into a “joint venture” between fewer, but bigger participants?

Are there questions I’m missing? Let me know: anthony.malakian@infopro-digital.com.



How tech evolution has changed the world of finance

Neal Pawar, the former CTO of AQR and current COO of Qontigo, chats with Anthony Malakian about some of the major trends that are changing how asset managers interact with the vendor community, and how this shift mirrors the most significant evolutions in capital markets technology over the last decade.

I don't normally like starting a column by establishing a source's bona fides, but I think it's important to highlight Neal Pawar's career before getting into the meat of the article.

Pawar was AQR Capital Management's CTO. He spent more than a decade at DE Shaw & Co., during a time when the fund's assets under management grew from just shy of \$1 billion to \$45 billion. He was chief information officer of UBS Wealth Management. He was group chief information officer of Deutsche Bank. He's held other roles, and today he's COO of Qontigo, a risk, analytics, and index solutions vendor that was born out of the merger of Axioma and Deutsche Börse's Stoxx and Dax units.

His nearly 30 years of varied experience in the industry makes Pawar an expert in the field of financial technology. But more than that, his passion and zeal for the subject comes through every time we've spoken.

I recently caught up with Pawar because I saw he was named a member of Talos' advisory board. I was a bit surprised because crypto, digital assets and DLT were never topics of our conversations. As I think I've made clear in previous columns, I'm not a fan of crypto and DLT/blockchain.

But I do like to hear opinions that are contrary to my own, so I wanted to understand his reasoning for taking on this role (in addition to being COO of Qontigo, and serving as a member of CAIS's board). And through our conversation, I think that I gained a different perspective on the evolution of the buy side and how asset managers and hedge funds view the buy-vs.-build debate today.

But more importantly, I think this column is the closest I've come to explaining how the various tech evolutions we've seen since the financial crisis connect. This will sprawl, but I hope you stay with me.

A whole new world

So first it's important to acknowledge certain trends that we've been hammering on about—a lot. First is the inexorable push to the cloud for banks, asset managers, exchanges and vendors. On top of that, capital markets firms, especially end-users, are becoming increasingly comfortable with the idea of not just taking from the open-source community, but contributing. (And it should be noted that even before Pawar was at AQR, the hedge fund was on the cutting edge of open-source, as it released pandas into the world.)

Additionally, end-users and vendors are increasingly embracing APIs as the preferred way to connect to data sources. So it is that data is easier to acquire, store and analyze, but providing actionable insights into that data—whether for trading, risk, settlement, and/or cyber—is becoming increasingly important.

And it's easier to provide context if systems are interoperable and if AI is used to help in these efforts. (Yes, I'll discuss crypto and DLT, but first, you'll have to suffer through 1,000 or so words.)

Pawar says that just 10 or 15 years ago, large buy-side shops would choose to build—to some extent or another—things like their order management, settlement, and accounting systems. But because of these macro tech shifts we're seeing—in addition to increased regulatory reporting requirements, the rise of passive investing, and the pressure to

build more diverse portfolios—the buy/build discussion is changing.

“You're at the point now where even at the tip of the spear—front-office, portfolio construction—a lot of the components there you can buy off the shelf,” in addition to your middle- and back-office systems, he says. “But these components do not come preassembled; you need to assemble them, and that's where a lot of the alpha and the nuances of what makes one fund or asset manager different from another starts to come in.”

“So if you follow the argument that everybody's moving to cloud over a timeframe at different paces, then ultimately you also have to come to the conclusion that people are going to be rebuilding their own rails in the process.” Neal Pawar, Qontigo

If you talk to non-engineers, you might credit this evolution to those three aforementioned market drivers: reg reporting, passive investing, and asset class diversification. But as any technologist who has experienced the nightmare of replacing a major legacy platform knows, it's the technology that is rapidly changing, and everything else follows as a result.

Over the last 10 years, the rise of cloud has led to the rise of software-as-a-service (SaaS) and managed services, and incumbent tech providers have had to reimagine their established but aging systems. Not long ago, Goldman Sachs tried to send a former employee to the slammer over some open-sourced code; today, GS could be considered one of the leading banks when it comes to open-source development. Additionally, APIs

are becoming increasingly prevalent as the preferred way for data and software providers to deliver their goods. And at the same time, advancements in the fields of machine learning and natural language processing are allowing quants to more efficiently identify trends and trading opportunities.

Recently, the Depository Trust & Clearing Corp. (DTCC) and consultancy Celent put out a report that looked to benchmark just how far along

in comparison to the FAANGs, which have seemingly taken over most every aspect of everyday life.

But Pawar believes that the world of finance is starting to catch up: cloud architecting; developing high-speed analytics platforms; using machine learning and NLP to create previously unseen insights; embracing open-source tools ... for an engineer coming out of university, it's an exciting field (and, again, I haven't gotten to crypto/DLT yet). On

"I think that's the reason you see a number of experienced technologists moving over to the vendor side. I think we all see that that's the direction of travel. And there are still some huge opportunities out there for vendors to emerge and be influential in some of the spaces that aren't quite as commoditized as, say, accounting systems. No one builds accounting systems today, but you still might find firms building optimizers or building portfolio construction things in the quant space themselves," he says.

"Unless that is your true alpha and the thing that really makes you different from the others, then at some point everyone is asking the question: What parts can I buy off the shelf and integrate, and what parts can I do myself?" he says.

So the key takeaway is that the buy-vs.-build debate is changing because the technology is changing. Regulators will bring forth new requirements, market structure changes will be made, and investment strategies will evolve ... but these things can only be answered for if technology evolves, first.

“

"One of the things that would always frustrate me 10, 15, 20 years ago, was: let's say you were a fund that traded one asset class, so you're an equities shop and you wanted to start trading fixed income. Oftentimes, to do that you would buy an entirely new vertical of systems and you'd have to put them together for that asset class." Neal Pawar, Qontigo

buy- and sell-side firms are in their cloud migration journeys. As you might expect, the ranges differ, but the cloud is where everyone wants to be.

Here's how Pawar sees it: Some trading firms, exchanges and vendors are already fully in the cloud; some are getting there; some are just getting started. As a result, you're seeing more deals being struck between the likes of Amazon Web Services (with, for example, Goldman Sachs), Microsoft Azure (HSBC), Google Cloud (CME Group), IBM (Bank of America) and I'll throw Snowflake in there (BlackRock).

"So if you follow the argument that everybody's moving to cloud over a timeframe at different paces, then ultimately you also have to come to the conclusion that people are going to be rebuilding their own rails in the process," says Pawar, who perhaps enjoys talking about the rails and pipes of the industry—the infrastructure and backbone of financial services—as much as discussing cutting-edge tech.

Pawar was drawn to the capital markets after graduating from Brown University with a degree in computer science because it was an industry that was developing cutting-edge analytics systems and ways to manipulate data. These are my words, but perhaps finance started to stagnate in the 2000s, especially

top of that, while quant shops like AQR, Renaissance and Jane Street, or banks like Goldman and JP Morgan will continue to attract top engineering talent, the vendor community around financial services is becoming deeper and more sophisticated, and thus more enticing to graduates with computer science degrees.

Sure, Pawar might be biased, but he was drawn to Qontigo because of the freedom to develop new products and work on interesting projects with a range of companies. Pawar is not alone.

This is just anecdotal, but since the pandemic began, I've seen several big-name executives leave end-user firms and go to a vendor: Bill Murphy (Blackstone CTO, now at Cresting Wave); Suvrat Bansal (CDO at UBS Asset Management, now at Clarista); Joe Lodato (Guggenheim Partners chief compliance officer and CISO, now at Capital.com); Anjala Gupta Reddi (chief data governance officer, now at Dow Jones); Nancy Selph (COO and CTO UBS, now at Avellino Lab); and Keith Lubell (CTO at Berkery Noyes, now at Meta). And those last two, you might notice, are examples of industry professionals taking their capital markets skills and applying them to other sectors, which is something of a mini-trend.

Pawar doesn't view these moves as entirely coincidental.

The crypto end of the equation

Again, though, the reason I reached out to Pawar in the first place was mainly because I saw he was joining Talos as an advisor. And if you're going to talk about evolution in the capital markets, the topic of crypto (and DLT ... but more on that in a minute) must be addressed.

First, something to know: Anton Katz, cofounder of Talos, reported to Pawar when the two worked at AQR. And Katz met his partnering cofounder, Ethan Feldman, prior to arriving at AQR, when the two worked at Broadway Technology. So Pawar has a direct connection to Katz and Feldman.

But the backgrounds of Katz and Feldman play into what Pawar has been talking about when it comes to pipes and rails. Katz, Talos' CEO, was head of trading technology at AQR. At Broadway, he held several roles, most recently as director of software. Feldman, Talos' CTO, spent a decade at Broadway and served as lead software engineer. Broadway, of course, is one of the leading fixed income trading platforms in the industry (which led to the whole ordeal of getting acquired by Ion Group, only

to have the company splintered off—and that's too complicated to get into for this column).

This is all to state that Katz and Feldman have hardcore backgrounds when it comes to building institutional-grade trading platforms.

Now, what has hindered crypto's ascent in the wholesale capital markets? Well, to name three issues, it's the lack of institutional-grade trading tools, clearing and market data. But major improvements have been made in just the last 12 to 24 months.

Pawar says that what he respects about Talos is the company is trying to take the lessons learned from AQR and Broadway, and apply them to the world of crypto to build the infrastructure that will not just underpin this nascent asset class, but that will allow it to more seamlessly be tied to more traditional asset classes, like equities, foreign exchange, and fixed income.

"One of the things that would always frustrate me 10, 15, 20 years ago, was: Let's say you were a fund that traded one asset class, so you're an equities shop and you wanted to start trading fixed income. Oftentimes, to do that you would buy an entirely new vertical of systems and you'd have to put them together for that asset class," Pawar says.

"What you really want to do is take the set of systems you have and add an additional asset class to them so you don't have to clone all that functionality. But that requires the other side of that to be mature enough to connect in. Previously, that didn't really exist in crypto, which is why I think a lot of institutional investors were not the first out of the gates trading here—it was much more prevalent on the retail side," he says. "I think now that institutions see crypto as an asset class that's not going away—and it will become more and more part of their end-clients' asset allocations or investment thesis—they have to have the rails and the pipes."

Talos, he says, is looking to become the institutional smart order routing system for crypto, powered through APIs that take care of a lot of downstream connectivity to the ever-expanding range of exchanges out there. This way, a trader or portfolio manager can look at it as just



another asset class to connect into. "So now, when you're plugging that into the trading systems and order management systems that the client might be running, they don't have to treat it any differently," he says.

(I'm not here to plug Talos' offering—this is simply to show how institutional trading tools are being brought to crypto and why Pawar thinks that's important.)

And, of course, blockchain

As we've written about extensively before, blockchain skeptics abound, and so many supposedly innovative startups and game-changing DLT projects have gone belly-up or disappeared into the ether.

Pawar acknowledges this: "So many of us on the banking side and the asset management side have been experimenting with DLT for a long time now—seven to eight years of prototypes and proofs of concepts. When you look at actual production you say, 'Well, how many of those are actually driving the financial system?' The answer is still pretty small."

With that said, he believes that DLT development will be important for the future evolution of the capital markets. First, crypto isn't going anywhere—it's a small asset class, but an asset class that's growing nonetheless. And being that crypto is a digital asset, and the fact that governments around the globe are looking at the expansion of digital assets

into the world of traditional finance, that means that the technology that underpins digital assets, DLT, is here to stay, as well. (We'll save whether DLT should be used to solve every back- and middle-office need in finance for another day.)

As I understand the situation, for Pawar, it all comes down to standards and standardization.

He says that with DLT, the technology itself is still a bit too new and there are still concerns about throughput. And critically, he says every bank and asset manager wants the payload—the format of the data structure that's living inside of the distributed ledger—to be as close to what their internal representations are to either own that IP, or to at least make the job of migrating to that system very easy for them. In this way, if it's an over-the-counter swap that's being traded on the ledger, a trading firm can have control to make "design choices" as to the terms and conditions of the swap and the structure of the legs of that swap, for example. True distributed ledgers are designed to be rigid.

Pawar likes open source in part because it encourages a set of standards, protocols and agreed-upon terms and definitions for data interchange. One needs to simply look at the Financial Information eXchange, or Fix, for proof as to why a protocol can be a boon for an industry. He says if a firm wants to switch out its accounting or even risk systems, that



journey can take upwards of two years. “But trade execution systems [are able to be migrated in] almost a few minutes in a routing table and it would work, and that’s because of Fix,” he says.

“For me, one of the things that’s missing in other parts of the financial ecosystem is we haven’t quite invented that ‘Fix equivalent’ in those areas. So there’s still a certain amount of stickiness around the choices that you make with your vendor landscape because you’re not at the point where you can say, ‘I’m going to get off of this platform and move to another one or try it out,’ without going through a pretty lengthy, often multi-year project of running [applications/data] in parallel, and mapping securities, trades, identifiers and what have you,” he says.

In the same vein, Pawar says the Fintech Open Source Foundation (Finos), has been “super helpful” at providing a framework—a checklist, if you will—for how firms can incorporate open-source tools into their organizations, understand what the liabilities and licensing needs are, and how to train people.

And as *WatersTechnology* first reported, Finos will host the common domain model for Isda, Isla and Icma. With the CDM, Pawar notes, Isda made “design choices” because of its standing in the derivatives space, so the industry is OK with that. DLT, again, is rigid and doesn’t quite allow for design choices.

“That’s one of the reasons [DLT] hasn’t taken off yet, but as that starts to take off, then you start recognizing that now, when I look at companies

like Talos, they’re used today for trading cryptocurrency because that is the primary asset class that lives on these distributed ledger-type rails,” Pawar says. “But as you start getting representation of all other asset classes in more of a digital asset representation, then there are going to be opportunities. Traditional trading systems might not interact with them in the same way, and you may want to use the rails that companies like Talos have built, because they’re more naturally geared for working with digital assets.”

So as digital assets mature and governments and regulators better understand how these can be incorporated into traditional finance, the nature of DLT can provide some semblance of standards and protocols that don’t currently exist today in other, say, OTC markets. From there, the vendors in the space can provide the infrastructure (rails/pipes) that will allow for firms to make those design choices, but within the structure provided by standards and protocols.

Obviously, this is not going to happen tomorrow, but it’s why for every DLT skeptic, there are three or four believers.

In conclusion

At the top, I told you this would be a sprawling trip, but let’s see if I can wrap this up nice and clean:

The cloud piece: Cloud led to firms being able to pull in more data (the alt data explosion), which led to more experimentation with more analytics

platforms (context is king), which led to more experimentation with AI (how to find correlations in a sea of information), which led to more experimentation with open-source and APIs (standardizing tools to better move/analyze data), which led to the rise of SaaS and managed services (thus the ease of coming to market), which led to firms wanting vendor systems to be more interoperable (well ... the interoperability push), and the need to simplify the engineering (low-code development, which I’m only mentioning now, but is becoming a very popular topic).

You can argue about the timeline, but those are the trends most aligned with development in the wholesale/institutional capital markets.

The crypto piece: Somewhat in parallel with the cloud piece, crypto was born (digital assets). While it’s been a chaotic space, efforts are being made to provide more institutional-grade stability (meaning tools and data). And underpinning digital assets are distributed ledgers (a technological form of standardization).

Now the cloud piece is something I’ve been forming in my head through years of interviews, and Pawar tried to help me to understand how these varied components fit together for a CTO/CIO at a bank or asset manager. The crypto piece is Pawar trying to help me understand why crypto/DLT will prove a major disrupter in the future—and I likely didn’t do a good job of fully articulating his thoughts.

So with all that said, I’ll leave you with this: I will have been covering capital markets technology for 13 years this October. When it comes to the most significant tech changes I’ve written about in that span, well ... this column encapsulates the industry as I’ve seen it evolve. Of course, there were regulatory pressures driving change (just look at the ever-growing field of regtech). There have been massive market structure shifts that have occurred and that are still underway. And the growth of passive investing has been massively disruptive.

But looking at the actual technologies being used to address these industry shifts—the above is the state of the union as I see it ... assuming I understood Pawar correctly. **WT**


**waters
tech**




Waters Wavelength Podcast

Out every week.

Tune in as we discuss the latest industry issues and speak to market leaders.



Subscribe at:
soundcloud.com/waterstechnology
or search the iTunes store for
“WatersTechnology”



Human Capital



OpenFin adds Sanders as chief digital officer

OpenFin, a desktop interoperability operating system provider, has appointed Vicky Sanders as chief digital officer. Sanders will drive digital transformation across the company's ecosystem of buy- and sell-side institutions, along with the global vendor community.

She most recently served as chief commercial officer for TP Icap's agency execution division, where she was responsible for overseeing the commercial strategy of the Liquidnet and Coex Partners businesses. Prior to that, she co-founded RSRCHXchange in 2014, a fintech operating as an aggregator and marketplace for institutional research, and led the business through the acquisition by Liquidnet. During her time at Liquidnet, she served as head of investment analytics, before Liquidnet's acquisition by TP Icap.

Pawar joins Talos as strategic adviser

Talos, a provider of institutional digital asset trading technology, has recruited Neal Pawar as the company's latest strategic adviser. Pawar is currently COO of Qontigo, a provider

of index, analytics, and risk solutions that is part of Deutsche Börse Group, and also serves as a board member at CAIS, an alternative investment platform.

He was previously managing director and CIO of UBS Wealth Management, group CIO at Deutsche Bank, and partner at both DE Shaw & Co. and AQR Capital.

Tradeweb appoints Aigrain and Yared as independent directors

Fixed-income trading venue Tradeweb Markets has announced that Jacques Aigrain and Rana Yared were joining its board as independent directors in August. Aigrain will chair the compensation committee and serve as a member of the nominating and governance committee. Yared will serve as a member of the audit committee.

Aigrain brings more than 30 years of financial services expertise and leadership experience to Tradeweb's board, including seven years as an advisor at Warburg Pincus and nine years at Swiss Re, where he was CEO from 2006 to 2009. Prior to Swiss Re, he spent 20 years in global leadership roles at JP Morgan Chase in New York, London and Paris. He currently serves as chairman of the board of LyondellBasell, chairman of Singular Bank SAU, and director of Clearwater Analytics. He previously held board positions at WPP, London Stock Exchange Group, LCH Clearnet Group, Lufthansa, Resolution, Swiss International Airlines, and the Qatar Financial Authority.

Yared is a general partner at Balderton Capital, a venture capital investor focused on European

founders. She joined Balderton in 2020 following 14 years at Goldman Sachs, where she was a partner in the principal strategic investments group and later in GS Growth. She currently holds board positions at the Penn Fund and Wabash College. This appointment marks her return to Tradeweb's board, having previously served as a director for five years prior to the company's 2019 initial public offering.

Klein to join Deutsche Börse as chief compliance officer

Marc Peter Klein will join Deutsche Börse as chief compliance officer at the beginning of 2023. He joins from Deutsche Bank after seven years as head of business line AFC corporate and investment bank Germany and EMEA. He spent several years in senior positions at Barclays and DekaBank.

Genesis Global names chief marketing officer

Genesis Global, a provider of low-code application development platforms, has named Jason Jhonson as chief marketing officer. Jhonson joins Genesis from Improbable, a metaverse technology company, where he served as CMO, and previously chief commercial officer. Before that, he served in marketing roles at GitHub, PayPal and Yahoo.

Tradeweb taps Pluta for next president


Tradeweb Markets, a global operator of electronic marketplaces for rates, credit, equities and money markets, has announced that markets veteran Thomas Pluta has been appointed its next president. He will join the firm as president-elect in October



Vicky Sanders



Neal Pawar



Rana Yared



and become president on January 1, 2023, when current president Billy Hult succeeds Lee Olesky as CEO. As previously announced, Olesky will retire as CEO at the end of 2022 and continue as chairman through 2023.

As president, Pluta will work closely with Hult and the rest of the executive management team to set and execute against the firm's strategic agenda, oversee the day-to-day business, and continue to drive a strong culture of collaborative innovation. He will report to Hult. A member of Tradeweb's board since 2017, Pluta will continue on the board of directors in his new role.

Pluta will join Tradeweb from JP Morgan, where he spent nearly 27 years and was most recently global head of linear rates trading and co-head of North America rates trading.

ACA Group announces new roles following Foreside merger

Financial services governance, risk, and compliance adviser ACA Group has hired Marie Luchet as managing director of ESG, heading up the practice's strategy and vision for Europe. The firm has also named Neeraj Karhade CFO.

Luchet was previously the director of continental Europe with Principles for Responsible Investment (UN-PRI), a global association in sustainable finance. Prior to this, she served as head of responsible investment with Ecofi Investissements and as an ESG sell-side analyst with CIC Market Solutions.

Karhade was previously the CFO of Transformco, a holding company for assets once held by the Sears Corp., where he was responsible for the company's financial functions,

EDM Council adds NY Fed, Goldman Sachs alum to board

EDM Council, a cross-industry trade association for data management and analytics, has appointed Linda Avery, Verizon's chief data and analytics officer, to its board of directors. Avery is the first to hold the CDAO position at Verizon.

Prior to joining Verizon, she was chief data officer and head of data and statistics for the Federal Reserve Bank of New York. Between July 2014 and August 2019, she was the architect of the organization's data strategy and developed capabilities in areas that included artificial intelligence



Linda Avery

and analytics, data visualization, and data quality. Prior to joining the Fed, she spent 20 years at Goldman Sachs, where she was a managing director in technology.



Thomas Pluta

including financial reporting, audit and compliance, treasury, tax, strategic planning, and corporate finance. Prior to this, Karhade was CFO of Silverline, a Salesforce partner in the health, financial services, and media industries.

LiquidityBook hires new EMEA business operations lead

LiquidityBook, a provider of cloud-native buy- and sell-side trading solutions, has hired Jonathan Cross as general manager for Emea. He will be based in the firm's London office and report directly to CEO Kevin Samuel.

Cross will manage all aspects of LiquidityBook's Emea business operations, including the continued development of the firm's European buy- and sell-side client base. He will play a key role in securing long-term growth opportunities internationally as LiquidityBook carries out its global expansion strategy. He will also oversee daily operations in the Emea region.

Cross most recently served as head of sales and business development at Messer Financial Software. He also served as COO at Broadridge Financial Services' asset management



Jonathan Cross

arm, and held the same role at Tradar, now part of SS&C Eze. Early in his career, he worked as a developer at several financial institutions, including Barclays Capital, JP Morgan, Landesbank Berlin, Barings, and Hoskyns Group.

Stepp steps in as FactSet CTO

FactSet has promoted Kate Stepp to serve as the data provider's CTO, bringing her engineering, product, and organizational knowledge to the role of leading the company's technology organization and overseeing its digital transformation strategy.

Stepp joined FactSet in 2008 and is currently senior director of product management within FactSet's research and advisory workflow solutions business.

Prior to this role, she was senior director of engineering within FactSet's research workflow solutions business. During Stepp's tenure at FactSet, she has been a key contributor to many strategic initiatives, including FactSet's web offering, cognitive-powered insights, and FactSet for CRM.

Stepp will continue to be based in London and report directly to FactSet's CEO, Phil Snow. [WT](#)

The power of FDC3 appD: A universal standard to distribute and discover applications

The latest evolution of FDC3 includes new specifications to its application directories. By Dan Schleifer, CEO of Cosaic, and Jim Bunting, global head of partnerships

If you have an iPhone, you get all your apps from the Apple App Store. It's convenient, but it's a closed, proprietary marketplace that controls what apps you can have, how much you'll pay, and charges app developers 30% for the privilege of selling their apps in the App Store.

In the world of finance, the future is about desktop interoperability, with in-house and third-party applications that seamlessly integrate to build streamlined workflows for users. The open FDC3 standard codifies how apps should communicate and no vendor has control over this. But how do you discover applications and add them to your smart desktop? The FDC3 model is the opposite of the App Store—it's an open standard for application directories called appD.

Enter FDC3 2.0—and its new specifications to appD. A game changer for the discoverability and organization of applications, universal adoption of appD has the potential to drastically improve the way workflows are built. It stands to benefit all segments of the capital markets: sell-side firms with large IT departments, slim asset managers and hedge funds that rely on vendor technology, middle- and back-office functions, and vendors themselves.

Additionally, it's important to note that appD differs from the proprietary content/app stores that have been the prevalent option for application discoverability.

The goal of appD is to create a universal mechanism to catalog FDC3-compliant applications and information about them. The appD server will contain information about every app available in an enterprise, reachable by all smart desktops in a company. By organizing application metadata (such as name, type, location, context types and intents it supports, etc.) in a consistent way, applications can be easily found, added to a

smart desktop, and used within integrated workflows—and, more importantly, surfaced to the user when they're contextually relevant to the work being done.

With appD, you can now see every application available for a particular workflow and can easily plug the best choice into your desktop environment.

Discovering the right tool for the job

Large firms often have hundreds of applications available across departments and functional areas, and no central place to find them. For users it can be a difficult or even impossible task to discover and evaluate the right tool for the job at hand, and to find newly-available applications. Even within a smaller firm, each team often has its own apps and doesn't know what else is out there. For example, the fixed-income team may not be aware of an app that the equities desk uses that might be useful to them.

AppD is a better way to discover, launch, and then (using your smart desktop) integrate the technologies that help you do your best work. If a user wants to find the right app to connect seamlessly to their workflow, that process transforms from a lengthy research and implementation process, to a task that requires just a few clicks.

Beyond internal application directories, large vendors with portfolios of applications—think FactSet, Intercontinental Exchange, Refinitiv and so on—will host their own application directory to show off all their wares. Your smart desktop will not only connect to your internal directory, but to these vendors' directories as well, so you can see the applications available to you (all, or based on entitlements of what your firm has licensed).

Finally, open directories will emerge—with the first likely to be hosted by the Fintech Open Source Foundation

(Finos), which hosts the FDC3 standards—that catalog all FDC3-compliant applications from smaller vendors and fintechs.

Because the applications showcased in all of these directories are FDC3-enabled, they can interoperate with each other immediately. They're already speaking the same language even if they've been developed by separate teams or separate firms, or built using different technologies—web versus .Net or Java, for example. From the user's perspective, it all just works.

An open-source app store alternative

A mission statement of FDC3 that is particularly relevant here is “to develop specific protocols and taxonomies to advance the ability of desktop applications in financial workflows to interoperate in a plug-and-play fashion, without prior bilateral agreements.”

The appD specification is an open-source, decentralized alternative to the current status quo. Currently, firms don't have a consistent way to combine applications from multiple providers into one directory for their users. For vendors, their offering can only be seen as a viable option for users if they are contained in a proprietary app or content store. With appD, we have an open system where there is no middleman determining what apps are available and to whom.

AppD is valuable to everyone because it allows for a cleaner, more straightforward competition between technology providers, with frictionless discovery, deployment, and integration of those applications. AppD shows the biggest, clearest picture of what's available: an open market that consists of all the applications available to improve the way you work—set free from proprietary boundaries. [Wt](#)

Sign up to *WatersTechnology* newsletters

- Tailor your newsletter experience to suit your areas of interest
- Receive exclusive content handpicked by our editors
- Update your newsletter preferences at any time



Visit www.waterstechnology.com/newsletters to update your newsletter preferences today

Empowering innovation



reddot winner 2022

Proven AI-enabled cloud solutions

SmartStream has been a trusted partner of the financial industry for more than four decades. We have always been at the forefront of developing transformative solutions for middle- and back-office operations.

Our solutions and services span the transaction lifecycle and incorporate the very latest AI and machine learning technologies. These can be accessed in a variety of ways - via APIs and micro-services, or deployed as managed services and cloud environments.

Find out why many of the top financial institutions globally continue to rely on our innovations.

info@smartstream-stp.com
smartstream-stp.com



Stand A18