

AUTOMATING OPERATIONS ON MISSION-CRITICAL DATA

*Why establishing best practices for a data-centric operating model
in the middle- and back-office matters more than ever*

greyspark
partners

SPONSORED BY

DU
CO

“ Financial institutions have to rethink the way they design not just their IT platforms but their business processes. Historically, the approach has been workflow-centric: functional capabilities are documented, embodied in pieces of software and chained one after the other. In that context, data is nothing but the sum of inputs and outputs of these systems.

Going forward, the approach must be data-centric: the common thread running across the entirety of the business is data. Systems and applications become mere components that happen to read, write or update the data at some discrete points along the trade lifecycle.”

Frederic Ponzo,

Managing Partner, GreySpark Partners

Post-trade operations departments in major financial institutions are beehives of activity, but unlike bees, middle- and back-office teams are not perfectly coordinated and aligned to one cause. Yet, those teams are critical to the operation of capital markets firms, both in terms of economic performance and risk exposure. In an environment with multiple trading partners, diverse data formats and interconnected software systems of different types and tenures, post-trade processing is non-negotiable.

However, professionals in the middle- and back-office can be seriously ill-equipped to carry out these important tasks. Operations workflows tend to rapidly multiply, grow in complexity and become riddled with idiosyncrasies. The results are pervasive duplication of effort, additional costs and, more importantly, increased risks of error.

Middle- and back-office personnel are typically organised into dedicated teams aligned to specific functional, asset-class and/or regional lines, and rely on a wide array of systems for both data sourcing and workflow execution. While the data they ingest, analyse and output again is used by multiple departments and even exchanged with external parties – such as market infrastructure providers, clients and regulators – the technology to manage it is often selected and deployed on a per-team basis. Spreadsheets proliferate wherever there are gaps in the technology stack, and processes that rely on the data on such splintered systems over time degenerate into siloed, labour-intensive and error-prone workflows that serve only to create cost, delay and risk.

Consequently, with an ever-increasing number of internal and external data hand-over points, many institutions find themselves operating a giant data management facility without central control and visibility. Skilled middle- and back-office professionals are spending their time clearing break notifications and waiting on their IT departments to fix technical blockages, instead of doing higher-value work. Even at this high cost, getting a clear view of the right data at the right time often remains elusive.

From Data Opacity to Data-Centricity

The target post-trade operating model for large financial enterprises focuses on data quality as a process enabler, rather than on exceptions as process disruptors. Automating the operations around data – removing the manual steps involved in capturing, structuring, enriching and validating data – is at the core of such a model.

This model demands that firms understand the types and sources of data that are needed for straight-through-processing and requires them to proactively seek to ensure the quality of that data. Exceptions that do arise must be used to identify and remedy data problems upstream instead of wasting valuable resources on cleaning up the downstream symptoms. The model requires a versatile data automation platform that can empower teams across the organisation to standardise and consolidate their processes, establish data quality best practices and monitor data quality trends. Such a platform is data format-, product class- and use case-agnostic, and is operated by end-users who benefit from controls at the user, role and process level.

From this data-centric foundation, the business can:

- Identify the most common causes of trade failure and take proactive steps to remediate them before they cause further problems;
- Find and review trusted data for tasks such as regulatory reporting and risk management;
- Respond faster to market events and changing business demands, such as regulatory changes, the introduction of new product lines and the onboarding of new clients, by immediately understanding the data implications of such a change, which empowers business users to build, adapt and refine automations;
- Simplify their application landscape, cut costs and improve productivity by removing and replacing legacy systems, alleviating the support burden on IT teams; and
- Cut operational risk by automating opaque and error-prone end-user-developed applications.

“Previously, we would get access to our data in one big dump, but now the team has complete control of it and can start work on it as it feeds out. We can meet our SLAs faster, and the Middle Office team can see what is being assigned to them in real-time ...”

Director of Global Optimisation,
large multinational insurer

Clearly, capital markets firms who embrace a data-centric operating model are at an advantage in an industry where most organisations are stuck using overly complex and expensive exception management models. This report examines the ways that organisations approach the development of a new operating model and explores best practices to achieve it.

Many Routes to Better Data

Each organisation must determine its own approach to managing data processes in the middle- and back-office, as the requirements, priorities and resources vary from firm to firm. Capital markets business and technology consultancy GreySpark Partners interviewed leaders and practitioners at buy- and sell-side institutions making progress toward a data-centric operating model.

GreySpark discovered that most firms begin by focusing on reconciliation. Reconciliation creates a significant amount of operational pain for firms, all too often forcing staff to spend the bulk of their day manipulating and checking data because its accuracy and reliability is poor. Firms that start by tackling reconciliation adopt one of two key transformation approaches: standardisation and consolidation. Both activities must eventually be performed, but which is tackled first depends on the operational needs of each organisation.

Figure 1. Approaches to data management technology strategy in the middle- and back-office



Standardising Custom Data Operations

Certain types of data, such as OTC derivatives trade data, are commonly processed either manually or semi-manually as those products are often bespoke to each client. The processing of this data is, therefore, slow, susceptible to human error and encumbered with a high incidence of breaks. Firms use customised, stop-gap versions of systems that were intended for another purpose, resulting in fragile and unintuitive data operations. Finally, where organisations have deployed a reconciliation platform that is either impossible or prohibitively expensive to customise, spreadsheets and other workarounds are used to bridge the gaps, which leads to increased operational risk.

Usually, these types of manual and makeshift operations are sustainable up to a point, but if data volumes grow past a certain level, or the risks become meaningful enough to be picked up by auditors, the tenability of the approach plummets. Solving problems with bespoke workflows requires a flexible data automation platform with a low- or no-code capability that allows firms to create custom processes themselves. This eliminates the need for workarounds, lowers risk and improves processing outcomes.

Standardisation in Action Tier 1 Global Investment Bank

Problem

Senior leaders at this bank were concerned by the huge costs and risk within the operations department. Processes failed to meet the quality standards required for the internal audit team and regulators, and the frequency of errors in the data was excessive. This was caused by a reconciliations operating model that involved the manual processing of 1,300 reconciliations either via spreadsheets or macros.

Approach taken

The firm's key priority was to reduce risk, and so the bank first identified the highest-risk reconciliations. These included data recs of transactions in the less-standardised asset classes such as derivatives and commodities. The firm then began to feed that data into Duco, a cloud-based data automation platform. All newly created reconciliation requirements are also handled by Duco's no-code platform.

Additionally, a manual ETL process that orchestrated the multitude of bespoke recs was replaced by Duco's Data Prep tool. Thus, the end-to-end lifecycle of each process, from the ingestion of data to its reconciliation, as well as the exceptions workflow, is standardised on a single platform.

Results

Business-critical data quality controls are now fully documented by Duco's platform, providing a clear audit trail. The processes have been de-risked thanks to the removal of forty end-user applications and 1,500 systemic data errors, with Duco automatically labelling 95% of reason codes and automatically allocating 100% of breaks for one team.

The operations department can respond much faster to the future needs of the business. Reconciliations that took months to build can now go live within two to three weeks, and are built by the operations team themselves, freeing up IT resources.

“Duco puts power back into the hands of the users. The ability to clean up your own mess is a game-changer.”

Head of Reconciliations

Consolidating High-volume Data Operations

A different approach is needed to solve problems that arise from the reconciliation of standardised but high-volume data, such as cash trades data. Here, multiple teams often need to perform similar tasks using the same data, but due to the siloed way the business is organised, this work is performed independently, and at high cost. Consequently, operations teams waste time managing reams of breaks, which causes costs to rise in line with data volumes. In core areas such as the trading of cash products, this struggle is exacerbated as systems are older and likely to be on-premises, with all the associated hardware and maintenance costs that that entails.

In environments such as this, the remedy is to reduce both the number of platforms and the number of processes, which can be done by consolidating operations onto a cloud-based vendor solution capable of handling high volumes of data. Significant savings can be made across software licences, hardware and IT, and reconciliation teams can focus their time on more high-value activities and collaborate on processes instead of duplicating work.

“Time-to-market used to be 2-3 months and it is now more like 2-3 weeks ...”

VP Strategy,
Tier 1 European Bank

“We needed a tool that would allow us to have all recs in the same place, and the application and all recs to have the same interfaces making everything easily viewable ...”

Operations Engineer,
Tier 2 European Bank

“The adoption of Duco was key to the standardisation, consolidation and harmonisation efforts in Operations. It is also very complementary to data quality.”

Head of Regulation and Recs IT

Consolidation In Action Tier 1 Global Investment Bank

Problem

The bank had an extremely fractured technology landscape for reconciliations, which led to high costs and made central control difficult to achieve. Disparate teams ran their own bespoke, but often duplicated, reconciliations platforms and there was an overall lack of transparency and standardisation. Many of the platforms in use were legacy systems running on-premises, laden with quirks and workarounds to accommodate functionality gaps. However, the reconciliation function was growing quickly and needed both a better operating model and the technology to support it.

Approach taken

The bank ran an internal review of the solutions already deployed and, in a test of the speed at which new reconciliations could be developed, Duco performed best. Consequently, the bank decided to deploy Duco as the single platform onto which all other reconciliations would be consolidated.

Results

The project set a new standard for process and technology streamlining. Consolidation saved approximately 2.5 million Euros per year across hardware, software and staffing, and managers in other business lines are looking to apply the same simplification ethos to their functions to unlock cost and efficiency savings across the bank. There was also an increase in the match rate from 60%-70% to more than 90% using Duco.

This bank's experience illustrates how data automation can be implemented in a controlled, limited way that quickly demonstrates value, followed by wider deployment across the business once the benefits are proven.

“Duco has helped us to get a lot closer to the mark in data quality and to see flaws in upcoming data before they cause issues ...”

Global Head of Digital Strategy,
Tier 1 Custodian Bank

Data-centricity: Thinking about the Data That Drives the Process

When firms standardise and consolidate their reconciliations, they gradually gain sufficient visibility over breaks to perform root cause analysis and can then set about fixing the causes of exceptions upstream. However, by thinking data-centrally and with the target operating model in mind, they can implement change that will forever eliminate the need for more sophisticated reconciliation controls downstream, because data problems are spotted and corrected upstream.

The key to this more proactive approach is the deployment across the enterprise of a fit-for-purpose data automation platform. Such systems combine data preparation, validation, reconciliation and reporting, and are accessible by teams across the business at different stages in the data lifecycle. For instance, a data preparation module could sit between front- and middle-office systems, owned and managed by a Head of Middle Office or a Chief Operations Officer to ensure that data is normalised, transformed, enriched and validated before it flows into downstream systems.

Such a platform is also able to integrate with the existing technology stack, ingest data in myriad formats, give users the ability to build data controls without code and provide advanced workflow and reporting. This can enable huge cost-efficiencies and reduce data handover points between systems. Often seen as part of digital transformation projects, deploying a data automation platform fosters operational agility, business user empowerment and continuous innovation.

From this starting point, financial institutions can rapidly reap quantifiable benefits: de-risking their business, achieving a quantum leap in productivity and federating post-trade processes around data they can finally trust. This constitutes a marked improvement from the current exceptions management-focused operating model, where significant effort and money are expended waiting for things to inevitably go wrong and applying a never-ending series of short-term fixes.

“Our growth trajectory is steep, and we were on the lookout for an operating platform that could cope with our need to constantly refine and automate our processes. Duco has been successful at that and delivered what we hoped it would.”

Global Head of Operations

From Reconciliation to Data Automation Tier 1 Global Hedge Fund

Problem

The hedge fund had grown by acquisition and the reconciliations technology stack was split across three systems and numerous spreadsheets. Maintaining duplicate systems was expensive and the reliance on spreadsheets for reconciliations was creating risk and opacity. The hedge fund wanted to establish a one-system strategy using a cloud-based solution that would complement the abilities of its middle-office staff by providing a high level of technical functionality.

Approach taken

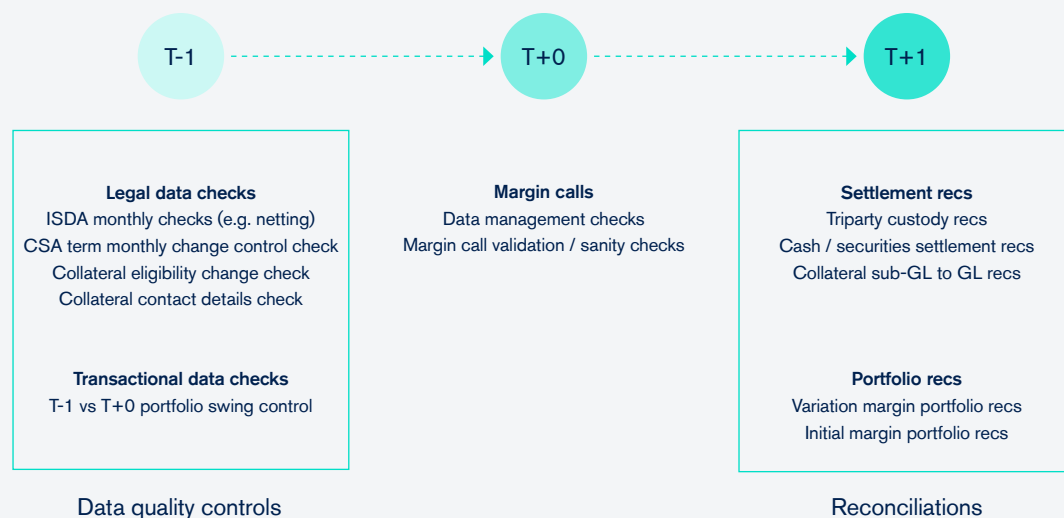
At first, the hedge fund deployed Duco's data automation solution across the organisation to replace the three different reconciliations platforms covering 1) collateral management, MTM & EMIR, 2) securities reconciliations for FX, equities and CFDs, and 3) cash reconciliations. Once the teams across these three areas were using the tool and had started to take advantage of the data preparation capabilities of the platform, additional workflows that could be automated using Duco's solution were identified. These included data aggregation and normalisation, regulatory reporting control and reference data. The interconnectivity of the platform enabled them to achieve greater levels of data automation since, as well as feeding data downstream via API, the platform was integrated with Python functions and robotic process automation tools.

Results

The hedge fund made considerable savings by consolidating reconciliation systems and decommissioning legacy hardware and technology infrastructure. The number of staff required to format data and clear breaks was reduced and the workforce re-focused to perform more advanced tasks, including creating new processes using Duco's no-code capabilities without help from the IT team. The firm has also experienced 50% AuM growth with no change in reconciliation headcount.

Today, the platform forms a core part of the firm's control framework and is seen primarily as a data integrity solution. Using Duco as a flexible toolbox to introduce more checks and controls into the data quality framework, as well as to improve daily workflows, has transformed both the efficiency and agility of the firm.

Figure 2. Examples of upstream data quality controls in collateral operations



Data Automation: a Roadmap to Data Centricity

Whether firms start with reconciliation and work their way upstream (as in the case study on page 9), or head straight to the source of 'bad' data, the journey towards a data-centric operating model is broadly similar. There is no one-size-fits-all approach, but there is a general adoption path that gives firms the best chance to minimise risk, cost and time-to-value. This is important because many firms have had negative experiences with technology change in the past, where they have been left with only higher costs and inefficiency.

By leveraging a cloud-based data automation platform, firms can start their journey to data-centricity with a chosen, relevant use case, keeping costs and risks low. This could be reconciliation-focused, such as eliminating manual processes for existing OTC derivatives reconciliations, or data-centric, such as cleansing and validating trade data on T+0.

The platform should integrate around the existing technology stack and be easily scalable – customers pay for what they need, knowing usage can be flexible. This is very different to an on-premise deployment, which can take months, often involves ripping out a previous solution and setting up complicated ETL processes to integrate with other systems, as well as locking the firm into an expensive and inflexible contract.

By targeting a single use case, teams can deliver return-on-investment quickly and demonstrate the positive impact of data automation to the rest of the business. As the case study boxes have shown, this tends to lead to interest from other teams and departments. Typically, the next step is to either expand the solution to related use cases in the same function or the same use case across multiple functions.

The final step, once the value of data automation has been proven, is to move towards deployment across the enterprise to allow the entire organisation to benefit from the data-centric operating model.

“We built new use cases on it and pushed the envelope to start developing it into a multi-purpose tool. We find we have been able to think more broadly with it and new use cases have emerged.”

Global Head of Digital Strategy

Data-centricity in Action **Tier 1 Global Custodian**

Problem

The custodian bank had wanted to resolve operational issues with a digital strategy that involved the acquisition and deployment of low-code / no-code tools. Reconciliations were managed on five legacy platforms, which the bank wanted to decommission in favour of one central platform.

Approach taken

Duco was selected for its no-code capabilities as well as its robust reconciliation functionality delivered through a modern technology architecture. The bank saw that operational risk was reduced as operations staff were able to use Duco's toolkit to develop processes without the involvement of IT, which was a large part of the appeal.

Results

Using Duco's platform, the bank's 1,000 reconciliations processes were consolidated into 200, and the cost and productivity savings from this, as well as the decommissioning of legacy platforms, has been a big success. The speed of the solution has proved to be an important benefit as new processes can be created in as little as two hours, and the time to build reconciliations has reduced from three or four weeks to days, overall. The firm can now build new processes for clients on the fly and rapidly report to business partners.

Importantly, Duco's data preparation functionality means that normalising data is easy, and any issues with data are spotted before the data has even been processed. The firm has started to see opportunities for data automation outside of reconciliations that were not previously possible. As the tool is very easy to use, it has become multi-purpose and attracted attention from other parts of the business that would also like to use it for new functions, such as parallel testing when moving data between platforms. Another unexpected use case is KYC, where the Duco platform has been used to evidence that client data is periodically reviewed.

True Data-Centricity

Data operations in the capital markets are on the cusp of coming of age – data management and automation technology is no longer a futuristic concept, but a reality in most large institutions. However, these technologies have usually not penetrated all areas of many firm's operations and, therefore, they are not being used as effectively as they could be. The middle- and back-office are the losers in the data automation race, but – as organisations re-evaluate their now-ageing legacy solutions – they are recognising the importance of including post-trade processing in any technology budget dedicated to data management evolution.

Data automation in the middle- and back-office should not be a luxury only available to a select few and it should not be deployed on a team-by-team basis. Instead, there should be an enterprise-wide data-centric operating model, enabled by a single data automation system for the entire post-trade landscape. This will save firms time and money both at the point of use and by ensuring that data entering any workflow is clean, reliable and auditable.

When addressing the core issues of reconciliation and operational controls along the trade lifecycle, the opportunity to solidify the foundation on which data operations rest should not be missed. Selecting a technology vendor that can support a forward-thinking data-centric strategy is an investment in improved productivity and reduced risk, and will deliver the agility needed for the post-trade landscape to thrive.

“ We viewed the system as a toolbox, and gave it to tech-savvy users to apply it to several business challenges that we had had for a long time. We have built it into processes, linked it to Python and RPA processes, and now it has become a central part of the toolkit available to my team to help them with processes that are run on a daily basis ... ”

Global Head of Operations,
Tier 1 Hedge Fund

**Use of published reports**

You are granted a licence to view and utilise this report for your own use. This licence is provided subject to the following restrictions: you must not republish, reproduce, duplicate, copy, sell, or otherwise sublicense or exploit the report or any part thereof. Use of the report or any part thereof for a commercial purpose is prohibited. To edit or otherwise modify the report without our prior written consent is prohibited. You must not redistribute material from the report, except for content specifically and expressly made available for redistribution in the report; nor use the report for any purpose other than that for which it was originally supplied and if the purpose of the report was not disclosed then the report should be considered to be for information purposes only.

Limitation and exclusion of warranties and liability

Whilst we endeavour to ensure that the information in this report is correct we do not warrant the completeness or accuracy of the report. To the maximum extent permitted by applicable law we exclude all representations, warranties and conditions relating to the report (including, without limitation, any warranties implied by law of satisfactory quality, fitness for purpose and/or the use of reasonable care and skill). We will not be liable for any loss or damage of any nature resulting from the use of this research report or any research products sold through it. We will not be liable to you in respect of any losses arising out of any event or events beyond our reasonable control. We will not be liable to you in respect of any business losses, including (without limitation) loss of or damage to profits, income, revenue, use, production, anticipated savings, business, contracts, commercial opportunities or goodwill.