Digital Regulatory Reporting

Pilot Phase 1 Report

Contents

How to navigate this document onscreen

returns you to the contents list

takes you to helpful Glossary

Disclaimer

The views expressed in the following paper are not representative of the views of the Financial Conduct Authority and Bank of England. They represent the individual views of DRR participants and so should not be taken as an indication of future regulatory direction.

1 Introduction

As regulators, the Financial Conduct Authority (FCA) and Bank of England (the Bank) need to collect information about the firms they supervise. These firms have to meet extensive regulatory reporting requirements. These provide information from a wide range of firms on a wide range of topics. But there is room for improvement in the way the current processes work. Regulatory reporting is expensive for firms and does not always provide regulators with consistent or high-quality data. It can also take time for firms to implement regulators' changes in collecting data on new risks or areas of interest.

The aims of the pilot

During 2018, the FCA and the Bank collaborated with Barclays, Credit Suisse, Lloyds, Nationwide, Natwest and Santander to carry out a 6-month pilot on 'Digital Regulatory Reporting' (DRR). The pilot explored how firms and regulators could use technology to make the current process of regulatory reporting more accurate, efficient and consistent. This included exploring the broader implications of those technological changes and developing a vision for what regulatory reporting might look like in the future.

During the pilot, the participants also sought 'to share the findings with the wider industry to facilitate feedback and allow for an industry effort to explore the feasibility of creating a new reporting mechanism'. This report aims to meet that objective by giving an overview of the pilot's work and findings.

The pilot's vision

The pilot participants developed a draft vision for regulatory reporting, based on developing solutions in 3 key areas:

i. standardising firms' data

ii. creating machine executable-code versions of regulatory instructions for how to create and deliver regulatory reports and

iii. developing a system to allow automated creation of regulatory reports which combine the standardised firm data and the machine-executable instructions from i) and ii) above

The pilot team also looked at options for how these solutions could be owned and governed.

What the pilot achieved

During the pilot, participants built a prototype using distributed ledger technology (DLT) to implement this vision for 2 cases. When used for live regulatory reporting, this prototype could potentially:

- improve consistency and data quality across firms,
- increase the efficiency of regulatory reporting and
- allow regulators to get data on new areas of interest more quickly

But there is a big difference between building prototypes for proof-of-concepts and delivering a system for use in real life cases. The pilot team started to assess which challenges would need to be addressed to successfully implement the DRR vision. This work included an assessment of the impact on the firms' and regulators' people, processes and technology. It also examined the rules that currently apply to regulatory reporting.

The pilot did more than provide new information about regulatory reporting. It also provided insight into how employees from two regulators and six firms can come together to try to achieve a common goal. This paper also shares those lessons.

Next Steps

The FCA, the Bank of England and a group of regulated institutions will be participating in a second phase of the pilot. This stage aims to identify which regulatory reports are appropriate for a DRR solution, whether there is value in investing in DRR, how best to create machine executable regulation and if and how to efficiently standardise firm data.

The paper is structured as follows:

- Section 2 sets out the context for this work, covering the current approach to regulatory reporting and other recent work in this area
- Section 3 gives more information about how the pilot was organised and its objectives
- Section 4 gives an overview of the DRR vision and prototype developed
- Section 5 sets out the pilot's key findings
- Section 6 sets out the next steps

2 Context

This section sets out context for this work, covering the current approach to regulatory reporting and other recent work in this area.

2.1 Regulatory reporting – current situation

Every regulated firm is required to submit data to regulators through regulatory reports. Regulation provides the instructions for how regulatory reports should be built and delivered. The amount of regulation and regulatory reporting has increased significantly in the decade since the financial crisis.

Given this increase in reporting and regulation, the complexity and time it takes firms to manage regulatory reporting has also grown. Firms provide regulators with reports defined in regulation, but regulators also make ad hoc data requests for data. Some firms have said this is challenging, with the amount and complexity of requests increasing over time.

The aggregate cost of this effort is significant. The European Commission's Fitness Check on Supervisory Reporting (EC's Fitness Check) estimated most firms' report-running costs to be around 1% of total operating costs. Industry feedback suggests that the cost of building or changing reports tends to be higher than running reports. This means the total burden on industry is likely to be higher than the figures quoted in the EC's Fitness Check.

There are many reasons why the process of supplying regulatory reports is now expensive. The process of building a regulatory report can be hard. The full set of instructions for compiling a report can be spread across many different pieces of interlinking regulation. The wording of the regulation might be insufficiently clear or difficult for firms to understand. Sometimes this reflects the challenge of writing a set of instructions that has to be understood and implemented by 50,000 firms. In other cases, firms need to make judgements which means it is difficult to create definitive, unambiguous requirements. International firms must meet requirements in several jurisdictions. They may have to repeat this process across many regulatory regimes, even when the regulatory reports contain similar sets of data.

Where possible, firms embed these instructions in their reporting systems, which reduces the ongoing cost of providing regulatory reports. However, some critical regulatory reports still involve many manual processes, making these reports particularly expensive to produce.

There can be serious legal and financial consequences for firms that get regulatory reporting wrong. To help manage this risk, many regulated firms buy external professional services to help them to interpret the regulatory requirements correctly.

Regulatory reporting also creates difficulties for regulators. Late and inconsistent regulatory data can damage regulators' ability to effectively supervise and monitor financial markets, identify harm and detect financial crime.

2.2 Regulatory reporting – other recent work

The pilot is the latest in a series of events run by UK regulators to look at regulatory reporting, as set out in Figure 1.

The first event was a TechSprint the FCA held in November 2016 on 'Unlocking Regulatory Reporting'. In November 2017, the FCA with the Bank of England held a follow up TechSprint on 'Model Driven Machine Executable Regulatory Reporting'. That event developed a proof of concept which demonstrated that a small set of reporting instructions could be converted into machine-executable code. Machines then used the code to automatically carry out (execute) the instructions, pulling the required information directly from a firm's systems. While the high-level process developed at the TechSprint demonstrated the approach was feasible, much progress was still needed for the concept to be considered proven.

Following an industry call for input, feedback statement, and a series of roundtables, the FCA and the Bank commissioned the 6-month DRR pilot.

3 How the pilot was organised

This section sets out the objectives and structure of the pilot, with some key lessons for organising future pilots.

3.1 Pilot objectives

The pilot was set up to explore how regulators and firms could use technology to make regulatory reporting more accurate, efficient and consistent. In particular, the pilot would explore how to make reporting rules and instructions less reliant on human interpretation and implementation, and so improve the quality of regulatory data.

The core objectives of the pilot, as set out in the Terms of Reference were to:

- develop a working prototype solution that demonstrates the end-to-end process for machine executable reporting.
- develop the prototype across two use cases – retail and wholesale – to help ensure broad coverage across the industry
- share the pilot's findings with the wider industry to ensure feedback and support industry effort to explore the feasibility of creating a new reporting mechanism.
- evaluate the potential costs and benefits of a new reporting mechanism compared to the current way regulators collect data from different sized firms

3.2 Pilot workstreams

The pilot's work was split across six workstreams. These aimed both to develop innovative technical ideas and to thoroughly assess the broader implications of altering the regulatory reporting process. The workstreams were:

- The Target Operating Model: to evaluate a potential implementation strategy for DRR. This included looking at future system ownership, the roles of firms and regulators under DRR, and data and technical standards.
- Technical Architecture: to evaluate the various architectural options for distributing code from regulator to firm, create machine-executable regulation and transfer data from firm to regulators.
- Data: to explore how to develop a data standard for financial data and allow regulators to automate regulatory reporting. This workstream also explored how to convert human (natural) language into a language that machines can understand and action (machine executable).
- Legal: to start considering the potential legal implications of altering the current regulatory reporting infrastructure.

- Policy: to highlight policy issues raised by the digitisation of reporting rules, identify and quantify the analysis and assessment resource needed to deliver machine-executable regulation for regulatory reporting and to evaluate options for future policy strategy, putting together a roadmap for the next phase of work.
- Product: to build an early version of a solution to demonstrate how the new system could be delivered technically.

3.3 Pilot use cases

The prototype built during the pilot implemented the DRR vision for regulatory reporting for two use cases. These use cases were firms' compliance reports with rules defined in UK and international regulation:

- **The Loan to Income (LTI) compliance check from PSD001:** This report is a quarterly view of the mortgages an authorised firm has sold. Regulators use the information to carry out checks to ensure firms are lending responsibly. PRA rules and FCA guidance (implementing a Financial Policy Committee recommendation) require that no more that 15% of mortgage contracts sold have credit exceeding 4.5 times the declared income (PRA Rulebook, CRR Firms, Housing Sections 1.11 and 2.1) and the LTI calculation checks this.
- **The compliance check for the Capital Equity Tier 1 (CET1) Ratio:** This is an updated measure of the minimum capital a firm must have against its assets. This ratio must be at least 4.5% of risk-weighted assets at all times (Basel III Standards, Part 1 Section I.A, Paragraph 50)

4 The DRR Vision for regulatory reporting

This section sets out a vision for regulatory reporting developed during the pilot and gives an overview of the technology used to implement the approach.

The pilot broke down regulatory reporting into 3 tightly linked but independent processes:

1. converting regulation into code (machine executable regulation)
2. defining standardised firm data
3. executing machine executable regulation against standardised firm data

While each process may need its own technology solution, those solutions must communicate seamlessly.

The pilot team used distributed ledger technology to quickly build a prototype that incorporated all three processes. This prototype implemented the process for regulatory reporting for the two pilot use cases: the CET1 and LTI compliance checks.

The team established virtual nodes to represent regulated firms and regulators. During the pilot, machine executable versions of the regulation were created as smart contracts. These smart contracts were loaded into the regulator's node in the system and distributed to the nodes of relevant firms (ie those to whom the regulatory rule applied). The codified regulation was then executed against synthetic data supplied by the firms in a standardised format. The results of the compliance checks were made available to the regulator and firm via a Graphical User Interface (GUI). Regulators had the functionality to schedule or run reports as required. For broader analytical purposes, the regulator was able to pull data directly from the firm node, using a smart contract executed as a request to the firm node via an Application Programming Interface (API), a tool for gaining access to data.

To show the system was suitable for firms of varying size, firm data could be uploaded to the system in an excel format. Firm nodes could also be deployed to a firm server remotely, allowing for firms to join the DRR system/platform quickly and easily.

5 Findings

The technology and processes used to build and run the prototype DRR system could not be deployed as a live, fully functioning DRR system without further development. Using a DRR system for live regulatory reporting would also require firms and regulators to make a number of non-technical changes, as well as changes to the rules for regulatory reporting.

During the pilot, these issues and changes were explored further. We present the findings of that work in the following 5 subsections:

1. converting regulation into code
2. a system for automating regulatory reporting
3. providing firm data to regulators
4. governing and operating the DRR system
5. implementing and adopting DRR

The final subsection looks at the lessons learned from operating a private/public pilot.

5.1 Converting regulation into code

Key Findings:

- even for rules that are intended to be clear and unambiguous, it is necessary to change how regulation itself is expressed to improve the efficiency of converting regulation into code
- there is no obvious solution to efficiently convert regulatory instructions expressed in natural language into code
- failing to identify the most efficient method does not prevent benefits from a DRR solution, but may make the DRR solution unsuitable for some regulatory reporting use cases

The pilot team gave significant consideration to which regulatory reporting requirements could be converted into code and the most efficient way to do this. They also examined the associated legal and governance implications from implementing DRR.

The prototype solution built during the pilot used two methods for converting regulatory logic into code:

- **Set parameters for regulatory content in a system**

1. *What is it?* Fixing certain logical relationships and concepts in code, while allowing the regulator or policy maker to easily change other relationships or data at relatively low cost.
2. *What did the pilot do?* Fixed the definition of a firm, a regulated activity and an obligation in the DRR prototype system – mimicking the structure of UK financial regulation. This allowed the regulator to easily assign regulated activities to firms and obligations to regulated activities.

- **Directly translating regulation into machine executable code.**
  1. *What is it?* Expressing the logic of regulation in code written in a general-purpose programming language like Java, C++ or Python.
  2. *What did the pilot do?* A combination of subject matter experts and engineers coded the logic of the regulation in JavaScript. The code was then executed as a smart contract on the DLT network (the distribution method chosen by pilot participants to share the regulatory logic with the industry).

Both mechanisms are currently used in regulatory reporting systems. However, the process of translating, agreeing and verifying regulatory instructions expressed in general purpose programming languages is difficult.

Therefore, during the pilot, the team identified 3 further options for converting regulatory content expressed in normal language into code. The team explored one of these options in detail during the pilot phase, but a lack of people and time prevented them from examining all the possible options in detail.

- **Rewrite regulation in a Domain Specific Language (DSL).**
  1. *What is it?* Regulation can be rewritten in a stripped-down, highly structured, machine-readable language – a DSL – that would sufficiently reproduce the structure and flow of regulation as understood by policy makers and legal professionals.
  2. *What did the pilot do?* The pilot team developed an early form of a regulatory DSL and were able to write small parts of regulation in this language. They were also able to turn regulation written in this language into machine executable code. However, the team did not have time to embed and run the generated code into the prototype system.

- **Leverage semantic technologies to output machine-readable artefacts.**
  1. *What is it?* Lawyers and policymakers write regulation in natural language that fully complies with the laws of English grammar. Further additional rules about how the regulation is drafted allow the output to be translated automatically into structured machine-readable content. This machine-readable content can be converted into code and mapped to the data held in firms' systems.
  2. *What did the pilot do?* The team did not have enough resource to explore this approach during this phase of the pilot. However, it has previously been tested for narrow use cases, for instance at the November 2017 TechSprint on 'Model Driven Machine Executable Regulatory Reporting'[1].

- **Generate the code using Natural Language Processing (NLP) technologies.**
  1. *What is it?* Technologies exist that allow for structured information to be extracted from natural language text. The structured information extracted by the NLP algorithm could potentially be used to populate a data model of a regulatory reporting system.
  2. *What did the pilot do?* The team had insufficient resources to explore this approach during this phase of the pilot. However, some firms currently use these technologies to extract key terms from legal documents.

1 During the TechSprint, participants built a system that linked the natural language definition of data point in a regulatory report (FSA001) to a mock-up of a firm's system. Regulators were then able to change the definition expressed in the regulation and change the data that was returned in the report.

The pilot identified several limitations on the process for producing machine executable code from regulation:

None of the options explored have yet proven they can be scaled up. The option commonly used today – human translation of regulation directly into machine executable code – is an expensive process at scale. However, the process may be significantly more efficient if it occurred once for all institutions, rather than separately at hundreds of institutions.

Even for instructions that are intended to be clear and unambiguous, it is necessary to **change how those instructions themselves are expressed to make it more efficient to convert them into an equivalent code format.** This is because many instructions written in natural language are not detailed enough to be translated into code. Rather, a human must make further assumptions to bridge the gap between natural language and machine code.

Both the content of the regulation and how it is presented create challenges. Instructions for regulatory reports are often embedded across many legal documents and published in pdf form. These documents are designed to capture legal and policy intent – not to be quickly and easily converted into code. It is also difficult to ensure the code in reporting systems accurately reflect the instructions embedded across many documents. This is particularly the case when these documents are subject to change and may be published by multiple regulators.

There is no obvious answer to what is the most efficient process to generate code from natural language regulation. Some solutions, such as using NLP technologies or certain semantic technologies, leave the process for writing regulations unchanged. There are obvious benefits here for regulators, but the risk of problems in translating between the language of humans and the language of machines remains. Failing to identify the most efficient method does not prevent benefits from a DRR solution. For instance, creating a standardised format to describe financial data would unlock efficiency gains in regulatory reporting. However, it may make the DRR solution unsuitable for some regulatory reporting use cases.

5.2 A system for automating regulatory reports

Key Findings:

- more work needs to be done to build a production quality DRR system ready for real world use.
- the prototype system can deliver significant efficiencies by centralising processes that firms currently carry out locally. It also demonstrated the possibility for real time regulatory reporting and a potentially dramatic reduction in the time and expense of regulatory reporting change.
- the system was able to reduce duplicate data storage and may potentially reduce data security risks through reducing data transfers between firms and regulators.

The prototype system realised significant efficiencies by centralising processes that firms currently carry out locally. In particular, interpreting regulatory instructions and producing logic to generate regulatory reports was performed only once and distributed to all firms via the DLT network.

The system also showed the potential for data to be reused across multiple regulatory obligations. However, as previously noted, ensuring regulatory reporting is done correctly has major financial and legal implications for firms. So the consequences of outsourcing those processes to a single body could be significant.

The prototype system demonstrated the possibility for real time regulatory reporting and dramatically reducing the time and cost of regulatory reporting changes. Implementing such a system for live regulatory reporting may pose real challenges and so could only be used in an appropriate governance framework (see subsection 5.4).

The prototype system could reduce duplicate data storage (and associated costs) and may potentially reduce some data security risks by reducing data transfers between firms and regulators. The data used to generate the compliance reports remained with the firms. Only the compliance result itself was shared with the regulator. This minimised the amount of data being transferred across the network. By making data available on demand, regulators may also store less data in their own infrastructure, helping regulatory IT budgets. Finally, by executing code within firms' own infrastructure, the sensitivity of data is potentially reduced at source before being passed to the regulators. Given GDPR obligations, this may be an opportunity to reduce the risk of transferring certain data between firms and regulators.

The prototype system showed that distributed ledger technology could meet some of the requirements of a shared regulatory reporting system. Benefits of the distributed ledger architecture included the ability to provide a secure channel to send codified regulations to multiple firms, to provide a single source of truth for a shared set of facts and a consistent environment that ensured the code ran successfully.

Work needs to be done to build a production quality DRR system that could be used for live reporting. Key system requirements such as security, performance and quality assurance were not included in the pilot scope. They would need to be incorporated in the design for a live DRR system. This may result in a change in the technology used in the pilot.

5.3 Providing firms' data to regulators

Key Findings:

- for regulatory reporting to be automated, the instructions need to be provided as a code which references data provided by firms
- ensuring that the format used to standardise data can be reused across multiple regulatory reports is critical to the efficiency of the DRR solution
- to ensure data quality is high, definitions must be precise and well understood by regulators and firms.

For regulatory reporting to be automated, not only do the instructions need to be provided as code, but that code ultimately needs to reference data provided by firms. To do this efficiently that data must be provided in a standardised format.

Ensuring that the format used to standardise data can be reused across multiple regulatory reports is critical to the efficiency of the DRR solution. Regular changes to the format or the need to maintain a large number of different formats for different regulations and/or regulators will significantly reduce or completely cancel out the potential cost efficiencies of regulatory reporting under DRR. This is one of the key problems of today's system. Firms need to provide data in a large number of different formats that are defined for each regulatory reporting requirement.

During the pilot, the team standardised mortgage data in accordance with the FCA's 'PSD001' schema, a transactional point of sale mortgage report. All data used during the pilot were test data stored in dummy databases run in the cloud.

The team identified at an early stage that standardising data according to a regulatory specific format was not a scalable solution. So the pilot's data workstream looked at the requirements and process for building a format that could be extended across multiple regulatory reports.

From a technical perspective, formats used to standardise and describe data so that it can be machine- accessible is called a data model. Data models are a number of logical boxes (or data points) with affixed labels. An associated definition provides the instructions for what data goes in the box. For instance, a data model may contain a box with a label 'entity name', with an associated definition 'The official legal name for an entity'. A system can then assure these boxes have been populated correctly by examining and validating known relationships between data in different boxes.

To ensure data quality is high, definitions must be precise and well understood by regulators and firms.

The data workstream focused on developing a method to improve the quality of definitions. This process was based on three key ideas:

- That issues in the understanding of definitions can be categorised into four fundamental problems[2].
- That definitions of data can always be defined in terms of other data. This creates a recursive pattern between definitions. This means that theoretically no definition can be completely unambiguous. However, for practical purposes, we can often express a definition in terms that are ultimately well understood.
- That those definitions can be expressed more precisely in a highly structured machine-readable format, potentially removing the need for human interpretation.

2 Definitional inclusion errors: a definition includes a set of things that were intended to be excluded in scope of the definition. Definitional exclusion errors: a definition excludes a set of things that were intended to be included in the scope of the definition. Definitional base errors: the terms used in the definition are not well understood and are themselves not defined. Definitional knowledge errors: the terms in the definition are understood but the firm does not have the information to understand whether their data meets that definition.

5.4 Governing and operating the DRR ecosystem

Key Findings:

- technical standards must be agreed on how the coded regulation will communicate with the system for producing the regulatory report.
- data standards will be required for how firms standardise their data before providing it to the system.
- under the new DRR vision, regulatory data could be available at much higher frequencies.
- while the pilot built a prototype system, a production system would require significant initial and ongoing investment. How that investment is paid for needs to be decided.
- any future DRR framework must ensure firms have the ability to verify data before it is published.

The DRR target operating model and policy and legal workstreams looked at some of the implications of the DRR vision on the rules currently governing regulatory reporting. The team also looked at the current options for building and maintaining technical solutions collaboratively across industry and/or through private-public partnerships.

The DRR solution envisages that regulatory logic is translated into code once. This translation may be performed by a regulator, a consortium of industry participants or a combination of both. Under the current system, every firm is responsible for interpreting regulation and implementing its requirements. **Shifting the responsibility for converting regulation into code to a central body raises questions about the legal basis of the coded regulation and liability questions about who is responsible if there are errors in the conversion.** There are also aspects of regulation that are firm-or situation-specific and therefore a single interpretation may not be appropriate.

The DRR solution also requires agreement on a number of key standards. Technical standards must be agreed on how the coded regulation will communicate with the system for producing the regulatory report. Data standards will be required for how firms standardise their data before providing it to the system. Security standards must be agreed to prevent data being accessed or used inappropriately. All these standards must be owned and maintained. **For the DRR solution to be used beyond the UK jurisdiction, these standards require international collaboration and governance.** A number of precedents for the governance and ownership of these standards, including Open Banking and the global Legal Entity Identifier (LEI), were considered during the pilot.

Central to DRR is a system that consumes regulatory code and executes it against data supplied by firms. The team built a prototype system during the pilot. **A system that could be used for live regulatory reporting would require significant initial and ongoing investment. How that investment is paid for would need to be decided.** The system must be more than just technology. Processes for onboarding users, managing system operations, and systems' maintenance would also need to be agreed.

Core regulatory reporting operates under a strict set of rules and procedures. Reports are typically submitted quarterly or monthly in a batch process.

Changes to regular reports require significant governance including formal consultation processes. Ad-hoc reports can be requested without a consultation process, but stringent regulatory processes still apply. Under the new DRR vision, regulatory data could be available at much higher frequencies. The cost of changes to regulatory reports may be significantly lower. This suggests **a new governance framework for regulatory reporting may be appropriate.** However, any such framework must ensure firms have the ability to ensure the data is correct before providing it to regulators and that there are appropriate safeguards in place to ensure data requests are appropriate and proportionate.

5.5 Implementing and adopting DRR

Key Findings:

- there are implicit risks and unknowns in delivering and rolling out of DRR
- implementing DRR will require changes to the internal technology and processes used by firms and regulators, as well as to human resources
- any adoption approach must be fair to all industry participants – not just those that participated in the pilot and related events

DRR is an ambitious project. **There are implicit risks and unknowns in delivering and rolling out of DRR.** This suggests it is prudent to take an 'agile' phased approach to delivery, and to recognise that the path to DRR becoming a reality is uncertain. Managing which regulatory reports could or should be in scope is critically important. These uncertainties mean DRR must be rigorously and realistically tested before DRR is potentially implemented, including through comparisons with existing regulatory reporting processes. While this increases the total cost of delivering the DRR solution before benefits are realised, greater work to understand the challenges of delivering DRR is necessary.

Implementing DRR will require changes to the internal technology and processes used by firms and regulators, as well as to human resources. Moving to a more 'on-demand' data-pull model would require changes to the way firms manage their regulatory reporting systems. Depending on the final model for the digitisation of regulatory reporting, regulators and firms would need time to adjust to a new process for writing regulation and the associated change in skill sets.

The fact that firms' data are stored in different systems and formats means squeezing those data into a single standardised format is a major challenge. To make this process easier, the pilot team considered how firms could supply data according to a number of interrelated formats while ensuring it was still possible to automate regulation. Technical solutions like these, as well as taking a pragmatic approach to data quality and consistency, would be necessary to ensure the roll out of DRR was not prohibitively expensive.

Should the standardised format used by a DRR system be embedded in firms' internal processes more broadly, the benefits for firms and regulators would be significant. Regulatory reporting could become a process in a broader group of firm operational processes, rather than a separate process in its own right. This would have major benefits as improvements in data quality used for internal purposes would also improve regulatory data. This raises the question of whether adopting DRR should follow other industry initiatives to standardise operational data.

Any adoption approach must be fair to all industry participants – not just those that participated in the pilot and related events. This could also delay the roll out of DRR.

The spirit of DRR has been one of collaboration based on experimentation and learning between regulators and firms. Making DRR a reality probably requires regulatory mandate. This may create tension if the interests of firms and regulators diverge.

5.6 Operating a public/private pilot

Through using agile techniques and a multi-firm, cross-jurisdictional approach, the pilot team identified a number of lessons to inform future similar projects:

- Public/private partnerships rapidly increase the diffusion of knowledge, improve the ability to identify problems and can lead to mutually beneficial outcomes.
- The value of multi-party collaboration was evident throughout, with the diversity of thought and contribution accelerating the rate of progress and quality of the results.
- Implementation, experimentation and agile ways of working can lead to a better understanding of the problem statement and solution requirements than theoretical desk-based exercises. However, this approach may not be suited to all aspects of designing and building a robust and secure system suitable for real world use.
- A collaborative environment with diverse skillsets and multiple organisations creates a stimulating and accelerated learning environment for participants.

6 Conclusions and Next Steps

The pilot process has built a better understanding both of how DRR could potentially be delivered, and the potential challenges of a roll out of DRR for both regulators and firms.

The benefits of public/private collaboration in the financial services sector is clear from the six months of work on the pilot. Firms and regulators could innovate and test solutions at a rate that is not generally possible in projects run by a single institution. To this end, we will extend the model for a further phase, beginning in February 2019.

The goal of the next phase of work is to close the gaps this pilot identified:

- **Data** – The goal of this workstream is to understand if the data model can be extended across multiple groups of products
- **Machine-Executable Regulation** – The pilot phase did not rigorously test the approaches for efficiently developing machine executable logic. The team will focus their efforts on testing three options that were not examined during the first phase of work.
- **Planning** – This workstream will focus on undertaking a feasibility study of DRR implementation in the industry. The team will carry out an initial cost benefit analysis to sketch out the business case for delivering DRR.

Appendices
Appendix 1: Glossary

Application Programming Interface (API)	A set of protocols and tools for building software applications and gaining access to data.
Distributed Ledger Technology (DLT)	A digital system for recording the transaction of assets in which the transactions and their details are recorded in multiple places at the same time.
Graphical User Interface (GUI)	An interface containing elements for user interaction with a system, such as buttons, windows and controls.
General Data Protection Regulation (GDPR)	An EU regulation on data protection and privacy for all individuals within the EU and the EEA.
Legal Entity Identifier (LEI)	A system for the creation and management of unique identifiers for legal entities that engage in financial transactions.
Natural Language Processing (NLP)	A branch of artificial intelligence that helps computers understand and interpret human language.
Semantics	The analysis of the meanings of words and the relationships between them.
Smart Contracts	Self-executing contracts with the terms of the agreement between buyer and seller being directly written into lines of code.


DRR
DIGITAL
REGULATORY
REPORTING