R3 Reports

# Applications of Distributed Ledger Technology to Regulatory & Compliance Processes

Josh Stark

# Contents

R3 Research aims to deliver concise reports on DLT in business language for decision-makers and DLT hobbyists alike. The reports are written by experts in the space and are rooted in practical experience with the technology.

# Applications of Distributed Ledger Technology to Regulatory and Compliance Processes

Josh Stark

May 5, 2017

## Introduction

Rising regulatory and compliance costs have led financial institutions to seek solutions using new technologies. Popularly referred to as "RegTech", this new field has become a key component of financial institutions' efforts to comply with the expanding set of global financial regulations that followed the 2008 crisis.

**Dodd-Frank** In the United States, the Dodd-Frank Act introduced significant reforms to financial regulation, ranging from the introduction of new reporting requirements for specific financial instruments to the creation of new regulatory entities like the Bureau of Consumer Financial Protection.[1] Dodd-Frank had a particularly significant impact on the market for "over the counter" (OTC) derivatives contracts, by introducing extensive real-time reporting obligations and requiring that many of these trades to be cleared through central clearing organizations.[2] In Europe, the European Market Infrastructure Regulations (EMIR) introduced a similar set of reforms aimed at increasing transparency of OTC markets and reducing systemic risk. On average, banks have spent $25 million on compliance costs alone with Dodd-Frank and EMIR as of 2015.[3]

**Basel III** In the aftermath of the financial crisis, the Basel Committee on Banking Supervision (BCBS) developed the most recent version of the Basel Accords - known as Basel III - a voluntary framework of regulation aimed at improving banks' risk management.[4] Intended to be implemented in 2019, Basel III will strengthen banks' capital and liquidity requirements. Complying with Basel III will require financial institutions to gather large amounts of data from across their business in order to assess their risk, ensure that they are within the prescribed boundaries, and submit reports to regulators establishing that they are in compliance.

**KYC / AML** Financial institutions must comply with many other regulations, including "know your customer" (KYC) and anti-money laundering (AML) regulations that typically require banks to establish the legal identity of their counterparties. Banks must build mechanisms that meet these regulations into their processes, and maintain records proving their compliance in the case of an audit.

[1] Dodd-Frank Wall Street Reform and Consumer Protection Act, Pub. L. 111-203, H.R. 4173, Title X

[2] Dodd-Frank Wall Street Reform and Consumer Protection Act, Pub. L. 111-203, H.R. 4173, Title VII

[3] "The Rising Cost of Trade Reporting." Sapient Global Markets, 2015 `http://www.sapient.com/content/dam/sapient/sapientglobalmarkets/pdf/thought-leadership/Crossings_Spring2015_TradeReport.pdf`

[4] The Basel Accords are voluntary frameworks. However, they are frequently adopted into law by national regulators.

In each of the above, ensuring data integrity has emerged as a key concern. Financial institutions must be able to meet these demands to keep and maintain high-quality data in order to report data, analyze risk, and maintain auditable records. Regulators have highlighted this as a key concern, both in the BCBS's Principles for Effective Risk Data Aggregation and Risk Reporting[5] and the recent attestation requirements from the Financial Reporting Standards Board.[6]

In the background of these expanding obligations, financial institutions struggle with technical debt: the burden of relying on outdated technologies at the core of their financial infrastructure. Many large banks use antiquated IT systems that are costly to maintain and require an ever-expanding number of bespoke software solutions to compensate for their limitations. Over decades, banks have grown, merged, been acquired, built new service offerings, and adopted fragmented infrastructure across business lines. The result is a patchwork of systems, often incompatible, that are spread across multiple jurisdictions.

Into this context, "RegTech" has appeared with new solutions. Typically, these address a particularly complex or difficult process that financial institutions must undertake to comply with some regulation. For instance, these solutions may include capturing transaction data from a trading platform, tagging it with additional metadata, and processing the resulting dataset into a report or applications that help firms establish the identity of a counterparty more quickly. Other solutions fall into the category of "big data" or analytics that help firms manage large datasets that must be aggregated, analyzed and compiled to provide useful information.

More recently, blockchain or distributed ledger technology (DLT) has emerged as a key component of next-generation financial infrastructure. While most attention has been focused on the potential for DLT to transform the basic infrastructure underlying many of the world's financial markets, these changes will inevitably impact the way those markets are regulated. By offering a new architecture for financial services, DLT may ease compliance burdens by simply eliminating some of the complex, expensive processes that burden financial institutions. Not only can regulatory compliance be made more efficient, but the technology presents new opportunities for regulators to design better, smarter regulations that can promote efficient, safe markets while reducing costs to regulated entities.

This research paper describes how DLT can be used as "RegTech", offering both financial institutions and regulators new tools to facilitate regulatory oversight.

**In Part 1**, we introduce distributed ledgers and provide a general overview of how the technology can be used in financial services.

**In Part 2**, we consider the use of DLT in financial compliance. By identifying shared processes across many types of financial regulation, we can consider where DLT may offer improvements. We then consider how DLT can enable new forms of interaction between regulators and financial institutions. We conclude by examining the requirements for privacy and confidentiality for distributed ledgers used for regulatory compliance.

**In Part 3**, we examine a specific area of financial regulation - transaction reporting for OTC derivatives - and look in-depth at how Corda, R3's distributed ledger platform,[7] could be used to facilitate the necessary processes. Specifically, we consider the post-Dodd Frank Act regulations governing the reporting of interest rate swaps.[8]

# 1 Distributed Ledgers in Financial Services

Blockchain technology was first introduced by Bitcoin, the world's first cryptocurrency. Since Bitcoin's launch in 2009, many variations on the underlying design have been built for different use cases. Throughout this paper, I use the terms "distributed ledger" and "distributed ledger technology" to refer inclusively to both blockchains and other types of distributed ledgers that, while using cryptographically secure records and some form of consensus service, do not use a

[5]"Principles for Effective Risk Data Aggregation and Risk Reporting (BCBS 239)" *Basel Committee on Banking Supervision* - `http://www.bis.org/publ/bcbs239.pdf`

[6]"Navigating the Year Ahead: Banking Regulatory Outlook 2017." *Deloitte Centre for Regulatory Strategy Americas*, 2016, p. 13 - `https://www2.deloitte.com/content/dam/Deloitte/us/Documents/regulatory/us-banking-regulatory-outlook-2017.pdf`

[7]Corda is R3's distributed ledger platform, designed for use in financial services – see: `https://docs.corda.net/_static/corda-introductory-whitepaper.pdf`

[8]More specifically, we examine the reporting obligations under 17 CFR & 43 & 45.

blockchain data structure.[9]

Distributed ledgers allow multiple parties to jointly view and edit one shared, consistent record of information. Rather than maintain separate records of some important fact - a balance of cryptocurrency, the state of some program, or the state of a financial relationship between parties - a distributed ledger provides one single record of this fact that is jointly held and managed by multiple parties. Each holds a copy of the record, is able to view it, and may be able to update it, but each copy is kept in sync with all others through what is known as a consensus algorithm.[10]

A distributed ledger is maintained by a group of computers called "nodes". These computers form the network that supports the distributed ledger: they receive and share transactions that update the ledger with one another, enforce its basic rules, and each keep their own copy of the ledger. In this way, the ledger is "distributed" - there is no central party who is responsible for it, but rather a group of participants who jointly manage it. Some distributed ledgers have many thousands of nodes spread across the world, while others might only allow a handful to participate in a closed network.

Distributed ledgers are typically append-only, meaning that information on them cannot be edited or deleted, only updated. For instance, a blockchain contains each previous state of the ledger - each "block" - going back to the very beginning of the blockchain, providing an ideal audit trail to track the history of some information. Distributed ledgers like Corda similarly track every update to the state of a financial agreement between parties.

One way we can distinguish between different DLT solutions is to ask whether the ledger is "public" or "permissioned". Public blockchains like Bitcoin or Ethereum allow anyone to run a node and view the full history of the ledger. "Permissioned" distributed ledgers limit the ability to act as a node and access the ledger to known entities. Some distributed ledgers, like Corda, break even further from the original blockchain design by limiting access to information about specific transactions to a small set of entities. This paper focuses on permissioned distributed ledgers. In the near term, these are the technologies most likely to be used in the regulated financial markets we consider below.[11]

Blockchains and distributed ledgers are often described as immutable, though this term must be qualified. A public blockchain is immutable in the sense that there is a very high probabilistic guarantee that data recorded on a blockchain will not be changed. The economic incentives used to produce consensus make it extremely difficult for any party to edit the ledger improperly. However, a permissioned blockchain is more flexible, as it relies on trusted relationships between known entities. Errors and unwinds of a transaction can be processed through a subsequent transaction, if the entities with control over that record agree on the change.[12]

Let's consider a simple example of how two parties could use DLT to record an agreement. Imagine that two parties intend to enter into a simple financial agreement - a futures contract on the price of gold. They would record the terms of this contract on the distributed ledger itself. There would not be separate copies of a contract, but rather one contract to which both parties have access. As the contract moves through its lifecycle, the ledger keeps a history of every change and update, creating an authoritative audit trail of the entire transaction.

[9] The precise definition of what is and is not a "blockchain" is often controversial and considered symbolic of specific political opinions. The choice of terminology in this paper should not be read as a statement on the relative value of different technologies, but simply a matter of convenience and clarity for the reader.

[10] The method used to create consensus varies greatly across different distributed ledgers. Public blockchains like Bitcoin use a system of economic incentives that reliably allows a network of thousands of nodes to agree on each new "block" added to the ledger. Permissioned ledgers, in contrast, may rely on existing trust relationships between entities to coordinate consensus.

[11] It is worth noting that many permissioned DLT are implementations of public blockchains, like versions of Ethereum designed to meet the standards necessary for enterprise applications. Others, like Corda, have no public-DLT equivalent, and were designed expressly for permissioned use between financial institutions. Further, it is expected that many DLTs - public and permissioned - will be able to interoperate in the future.

[12] A distributed ledger where parties intend to occasionally alter historical data will need to design proper processes to ensure that this is only done when necessary (for instance, to comply with "right to be forgotten" laws for certain use cases). Because this would require each participant who holds a copy of the information to agree on a change, according to a process defined in advance, in principle this would not compromise the authority of the data from a regulator's point of view.

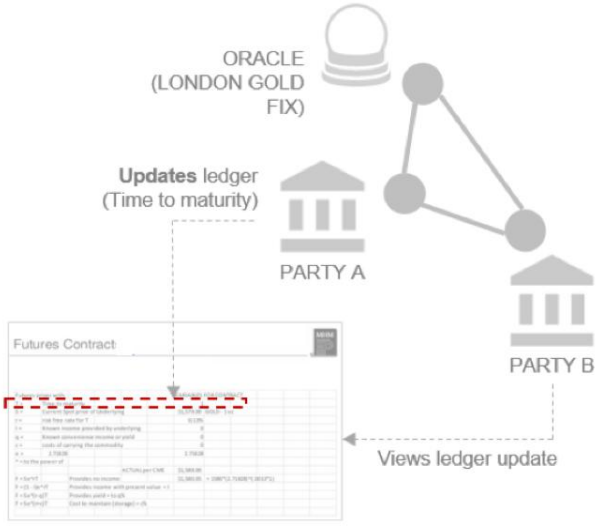


Figure 1: Multiple parties collaborate on a shared ledger used to record the state of a financial agreement

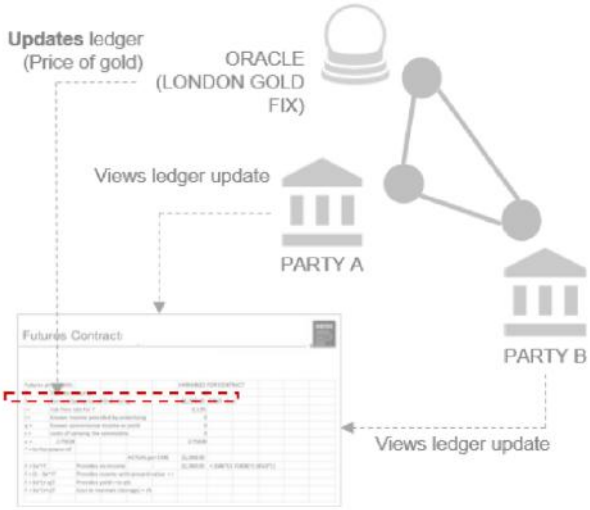


Figure 2: An **oracle** is used to obtain the price of gold, necessary to resolve the outcome of gold futures contract

The contract could be partly articulated in smart contract code, meaning that when it is triggered - for instance, by checking the price of gold on an agreed-upon date - it can execute automatically, resolving the outcome of the contract.

Many "smart legal contracts"[13] of this kind rely on a trusted source of outside information, known as an **oracle**. In the example above, an oracle provides the price of gold necessary for our futures contract to resolve. Many financial agreements depend in some way on information that must be obtained from a trusted source, like an interest rate, commodity price, or other market data.

## 2 Applying DLT to Regulatory Processes

Financial regulation takes many forms, and the processes required to comply with those regulations are likewise varied. However, we can observe that many regulatory processes share certain basic features. For instance, complying with any financial regulation requires the ability to se-

[13]See: J. Stark (2016). "Making Sense of Smart Contracts" - `http://www.coindesk.com/making-sense-smart-contracts/`

curely document and store information. The ability to share authoritative information securely with other entities - like regulators - is another fundamental requirement.

By considering these general processes, we can draw insights about how DLT can have applications in a variety of regulatory contexts. Broadly, regulatory compliance involves the following five general processes:

- **Recording and storing information:** the foundational requirement for regulatory compliance.
- **Aggregating data:** required where, for practical or legal reasons, data is stored in multiple systems or locations.
- **Performing operations on data:** data used in regulatory reporting must often be processed or analyzed in some way before being passed to a regulator. For instance, applying internal financial models to determine compliance with capital adequacy requirements.
- **Sharing information with other entities:** regulatory reporting obligations require that firms be able to share information securely with regulators, and occasionally with other entities.
- **Ensuring data integrity:** firms must employ processes to prevent or correct errors introduced by the processes described above.


Figure 3: The basic processes involved in regulatory compliance

## Mapping DLT onto Regulatory Processes

Considering how each of the above processes would be accomplished using DLT provides a clearer understanding of the technology's potential to address regulatory challenges.

In part 1, we described briefly how DLT can be used to facilitate financial transactions. Considering the regulatory processes described above, we can now see how DLT can facilitate simpler, more efficient regulatory compliance structures.

**1. Recording information.** Complying with financial regulations of any kind requires recording and securely storing data. For instance, firms rely on systems to "capture" transaction data from trading platforms and store it securely, where it can later be processed into the forms necessary for financial reporting. In other cases, firms may rely on manual processes to record information about financial events.

Firms are typically legally required to retain records of certain information about their clients and clients' business activities for a specific period of time. This includes, for example, FinCEN requirements to keep and maintain account opening data[14] or the obligation of firms to record and regularly publish financial information in their P&L.

The processes used to capture, record, and store this data can be simplified through the use of DLT. When a distributed ledger is used to perform a basic financial function - like track account balances, or the state of an agreement between parties - this information is stored in a single authoritative source: the distributed ledger.

Instead of creating a record of a financial fact - for instance, by capturing data from a trading platform and storing it in a separate database - we simply refer to the transactions themselves,

[14]BSA Electronic Filing Requirements for Reports of Foreign Bank and Financial Accounts (FinCEN Form 114) - `https://www.fincen.gov/sites/default/files/shared/FBAR%20Line%20Item%20Filing%20Instructions.pdf`

recorded on our distributed ledger. In order to report this information to a regulator, we simply give regulators access to the authoritative record of the transaction.[15]

This inversion of the typical process addresses many common problems with recording financial information. Often, the process of creating additional records is the source of errors.

For instance, under the European Market Infrastructure Regulations (EMIR), every derivative contract must be assigned a Unique Transaction Identifier (UTI).[16] Under EU law, both parties to the trade must report it to a data repository, and the UTI is used to match those reports to one another. However, this process is prone to errors. UTIs may be generated incorrectly, or not properly shared between parties. In 2014, the DTCC reported that it was only able to match 40% of the trade reports it received.[17] Similar issues exist with Legal Entity Identifiers (LEIs), which are used to clearly identify the counterparty of each trade. Under the upcoming MiFID II regulations, reporting entities will be required to obtain and include LEIs for every trade.[18] In the US, many financial regulations, such as CFTC and SEC swap data reporting mandate the use and reporting of LEIs.

Use of a distributed ledger might avoid the problem of matching entirely. There would not be separate records of a trade held by each party, but rather one single record shared between them. A UTI, LEI, or any other information contained within that trade would remain shared between every party that requires access. Quite simply, it would be impossible to "mismatch" reports, because there are no reports to match. There is only a single authoritative record, recorded on a distributed ledger. Practically speaking, this would be beneficial for reports such as FINRA TRACE reporting where regulators provide a report card on the number of mismatched and late fixed income trades reported.

More fundamentally, on a distributed ledger there would likely not be a need for UTIs in the first place, since they were introduced primarily as a method for trade reports to be matched - a process no longer necessary on a distributed ledger. If we did require a method to identify a particular transaction or record, we could simply use the unique hash of the data, rather than a UTI.

**2. Aggregating data.** Regulatory processes typically require aggregation of data from multiple sources. Many financial institutions operate multiple, sometimes incompatible, legacy IT systems that require complex aggregation processes to draw together the necessary data into a single record.[19] In other cases, the information required for a certain report to a regulator may come from many different business lines operating in different jurisdictions.

This is not simply a technical problem that can be solved through the use of a distributed ledger - there are practical and legal reasons that data will always be stored across multiple sources. For instance, there are requirements for "recovery and resolution" plans[20] that require systemically important banking functions to operate independently of each other, sometimes in distinct legal entities. Data localization and privacy laws may require certain information to remain within national borders, necessitating multiple data sources for multinational financial institutions.

Aggregation will still be necessary. However, the meaning of "aggregation" in this context would change. The process of "combining" data sources from multiple DLT platforms would instead be a process of collecting a series of links that point towards some set of records on a distributed ledger. Again, rather than create a separate repository of data that has been aggregated from other sources, we would instead have a comprehensive "live" view that shows us an aggregation

[15] The description here is generic to DLT. The way this is specifically implemented will depend on the platform. On Corda, access to information about specific transactions is defined by the "flow framework", which is designed by the developers of the Corda Application or CorDapp. If a flow framework is designed to facilitate regulatory reporting, then an authoritative record of the transaction will be pushed to a regulator node.

[16] Commission Delegated Regulation (EU) No 148/2013 - http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2013:052:0001:0010:en:PDF

[17] F Maxwell (2014). "Majority of Emir derivatives reports cannot be matched, say repositories" Risk.net - http://www.risk.net/regulation/emir/2335669/majority-emir-derivatives-reports-cannot-be-matched-say-repositories

[18] Regulation (EU) 600/2014 - http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.173.01.0084.01.ENG

[19] "RegTech in Financial Services: Technology Solutions for Compliance and Reporting." Institute of International Finance, 2016, p. 8 - https://www.iif.com/system/files/regtech_in_financial_services_-_solutions_for_compliance_and_reporting.pdf

[20] See, for instance, the Financial Stability Board's Key Attributes of Effective Resolution Regimes for Financial Institutions http://www.fsb.org/wp-content/uploads/r_141015.pdf, which is being implemented through the European Bank Recovery and Resolution Directive http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32014L0059

of different information stored in the authoritative ledger. This could be a powerful capability for compliance with BASEL III and similar regulations that require banks to understand their liquidity position across multiple products at a single point in time.

**3. Performing operations on data.** Regulators often require that firms process or analyze data, rather than simply record or report it. For instance, firms must apply internal models to determine whether they meet certain capital adequacy or other risk requirements. More simply, firms will generally have to prepare data into a report that meets some regulatory standard, transforming data to comply with differing legal definitions for key financial terms in different jurisdictions.

Similarly to point 2 above, DLT does not remove the need for financial institutions to be able to analyze data. Risk reporting will always require analysis of large data sets to determine the risk profile of a given financial institution and whether it is meeting a certain standard.

However, if distributed ledgers provide higher quality data with fewer errors, this could have implications for regulatory processes involving analysis. For instance, banks can fail Comprehensive Capital Analysis and Review (CCAR) due to data quality issues and the inability to trace that data to its source, even when they have sufficient liquidity to manage a stress scenario. If data quality were raised across the industry, regulators would be able to focus their efforts on conducting meaningful analysis drawn from the distributed ledger as "single source of truth" rather than spending time on understanding the individual systems of record at each financial institution. Higher confidence in the data underlying these models could be a basis for reducing capital requirements without necessarily increasing risk.

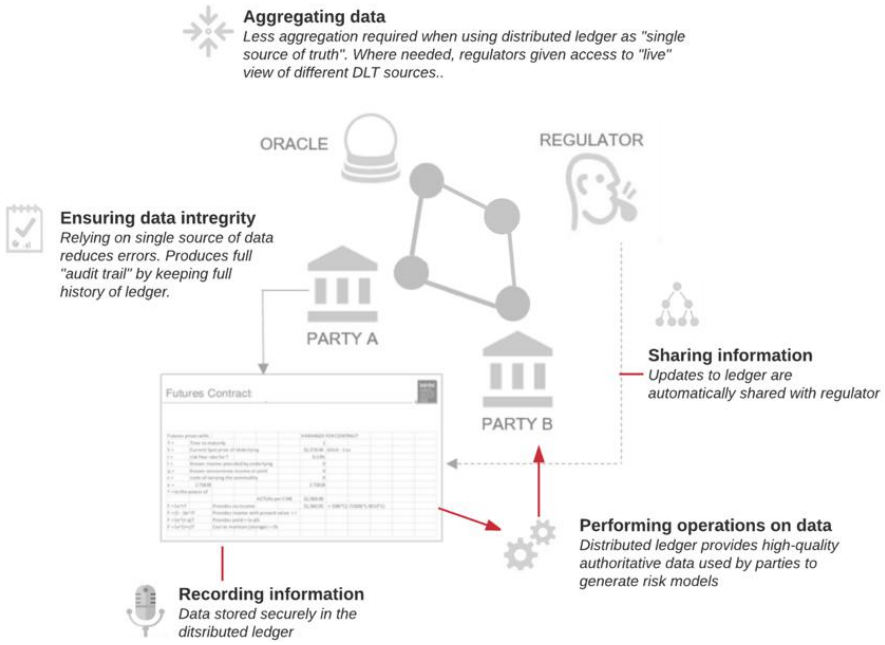


Figure 4: Basic regulatory processes applied to distributed ledger technology

**4. Sharing information.** Reporting obligations require that firms be able to share information securely with regulators.

Reporting obligations would be satisfied simply by giving regulators automated access to the authoritative record of a transaction, rather than create a separate reporting process as exists today.[21] In this way, reporting is built directly into the process of creating and executing financial agreements – as soon as a transaction occurs, the regulators see it happen. Since we are not creating additional, separate records of transactions, there is less need for reconciliation of data in the traditional sense.

Other components of reporting obligations could also be improved. Some forms of financial reporting have attestation requirements - some entity must "attest" to the validity of the data

[21]The description here is generic to DLT. The way this is specifically implemented will depend on the platform. On Corda, access to information about specific transactions is defined by the "flow framework", which is designed by the developers of the Corda Application or CorDapp. If a flow framework is designed to facilitate regulatory reporting, then the authoritative transaction will be pushed to a regulator node.

being reported. Wherever some outside party is needed to attest to some information inserted into the distributed ledger, an automated process could be created whereby that entity's cryptographic signature is obtained and attached to the appropriate information.[22]

**5. Ensuring data integrity.** Firms employ processes designed to ensure data integrity throughout all of the processes described above. For instance, reconciliation between different entities ensures that their records continue to match after being shared. Firms also generally create metadata which provides an auditable trail for particular records.

Many typical errors will be avoided simply through the use of a distributed ledger. When firms do not need to make multiple copies of some record and distribute it among many parties, there are fewer opportunities for errors to be introduced. The reliance on a single authoritative distributed ledger will reduce errors and the need for separate reconciliation, verification, and auditing processes.

Compliance processes require that financial institutions maintain authenticating metadata that can be used to attest to the provenance of some information and track its history from creation to reporting. Distributed ledgers achieve this by maintaining a full record of every update or state change in its history. A distributed ledger "automatically" produces an auditable trail and full history. Anyone provided with the correct permissions - for instance, a regulator - would be able to view the history of any particular financial agreement or transaction.

For instance, on Ethereum a regulator could view the history of a given contract-object by looking "back" to each block since the creation of that record. On Corda, a regulator would be given access to all dependencies and history of a state-object that represents a financial relationship.[23] This process is described in greater detail in Part 3.

## Opportunities for Innovation in Regulatory Technology

The previous section focused on how distributed ledgers can be used to facilitate and improve upon existing regulatory processes. However, DLT also introduces new possibilities for financial institutions and regulators. The new model for financial infrastructure offered by DLT may provide new tools to promote safer, more efficient markets.

With the generic DLT reporting structure described above, the marginal cost of reporting additional on-ledger[24] information to a regulator is very low. By default, a regulator who has been given access to a distributed ledger can have access to all of the information contained within a transaction.[25] We would have the inverse scenario to today: it is actually more complex to hide specific information within a given state-object (through transaction tear-offs[26], zero-knowledge proofs, or some other technique) than it is to simply reveal the full state of the ledger.

Not only does this reduce costs for financial institutions, but it might also enable regulators to demand more reporting than currently, while still lowering overall costs to the industry, since the marginal cost of additional on-ledger information is so low.

This same characteristic could allow firms to adapt more easily to changing regulatory requirements. Reacting to a new regulation that requires additional reporting would, in some cases, be as simple as "un-hiding" that data from the regulator where it was previously hidden.

Further, DLT may enable new methods to enforce regulatory compliance. The examples we have considered above have focused on how firms record, manage, share and verify data. Other regulations require firms to follow certain processes. "Know your customer" (KYC) and "Anti-money laundering" (AML) regulations are one example: when engaging in certain activities (like opening bank accounts, or transferring funds) firms must verify the identity of their customer or counterparty according to a legal standard.

With DLT, it would be possible to enforce these rules more directly. We could build rules into the platform itself, or into a specific application on top of the platform, such that certain

[22] We might also expect that some of these attestation requirements would become redundant. For instance, where they are required to provide greater confidence that a given report represents the real state of some entity or relationship. With a DLT, we would have a much higher confidence in this data in general, reducing the need for third-party attestation.

[23] Assuming that this information is within the regulator's jurisdiction.

[24] This is an important qualification. Reporting that requires additional off-ledger information would obviously still introduce marginal costs. For example, valuation data reporting may require firms to access market data and supply that to regulators. See Part 3 for a longer discussion of this point.

[25] The amount of information shared "by default" will depend on the specific DLT platform used. On Corda, the information that is shared between nodes is defined by the flow framework created for a specific application.

[26] See: `https://docs.corda.net/merkle-trees.html`

transactions could only complete once some condition is met.

For example, to comply with KYC standards, we could embed into the logic of a transaction a requirement that a counterparty provide sufficient documents such as driver license and proof of address to confirm its identify. The client would only complete onboarding once the bank received all KYC documentation and an oracle or a trusted third party attested the information provided is correct. A regulator could have greater confidence in this process, since they can see and verify for themselves that a given contract contains code that enforces the rule.[27] This is "regulation by design" - the mechanism used to carry out basic banking functions is constructed such that it must follow a certain rule, and the code that executes that rule is visible to the regulators.

## Privacy in Distributed Ledger Technologies

One fundamental challenge for distributed ledger technologies is preserving the privacy and confidentiality of participants and their data. Users of a distributed ledger - in particular financial institutions - may not want to share the details of every financial relationship with every participant in a distributed ledger, which may include their competitors. Further, in many cases there may be legal requirements that prevent financial institutions from sharing data, or storing it outside of certain jurisdictions. This emerges as a key design issue of a DLT for use in regulated industries and regulatory processes.

Permissioned ledgers can address this class of concern by limiting access to the data and relying on existing trust relationships between entities. In this way, they follow the model of existing financial infrastructure and are able to protect privacy and confidentiality in a manner familiar to financial institutions.

Public blockchains like Bitcoin or Ethereum do not limit access to their ledger. They are open in the sense that anyone can participate in the network by running a node, and they are transparent in that the details of all transactions are replicated to all nodes and visible on the ledger. At best, they offer a participant a pseudonymous identity, which can always be linked to a real-world identity by information leakage.

The European General Data Protection Regulation (GDPR)[28], which will come into effect in May of 2018, illustrates the privacy challenges for DLT. The GDPR places restrictions on how information about EU citizens may be used. In a forthcoming paper[29] written in collaboration with R3, Dr. Jana Moser argues that public blockchains like Ethereum would likely be subject to the GDPR.

Many of the obligations under the GDPR would be difficult to comply with for any entity using a public blockchain to store information pertaining to EU citizens. The "right to be forgotten"[30]allows an individual to demand the erasure of information under certain conditions, something that would be impossible in many cases with a blockchain like Bitcoin or Ethereum. Further, the GDPR imposes requirements on data that is transferred outside of the EU.[31] Using a public blockchain, it would not be possible for an entity to comply with the GDPR because the recipients (i.e. nodes) that hold the data outside of the EU, as well as their location, would be unknown.

Research into privacy and confidentiality protecting technology for blockchains is an active area of research and development. In a recent paper[32]published in collaboration with R3, Danny Yang, Jack Gavigan, and Zooko Wilcox-O'Hearn survey privacy and confidentiality-protecting technologies for blockchains. Techniques like ring signatures[33]and stealth addresses[34]can be used to protect the identity of a participant, by hiding the public key associated with a transaction. Others,

[27]Though note that the regulator must also have confidence in the process used to verify the documentation, which happens "off-chain" - a person must still view the documents, and then signal to our contract-code that they have been reviewed. Eventually, use of cryptographic identity may improve this kind of process – the transaction would only complete once it receives a cryptographic signature associated with a known legal entity.

[28]Regulation (EU) 2016/679 - `http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_2016.119.01.0001.01.ENG&toc=OJ:L:2016:119:TOC`

[29]J Moser. "Application and Impact of the European General Data Protection Regulation on Ethereum based Blockchains" (forthcoming)

[30]Codified in Art. 17 par. 2 of the GDPR

[31]See Articles 44-49 of the GDPR

[32]D Yang, J Gavigan, Z Wilcox-O'Hearn (2017). "Survey of Confidentiality and Privacy Preserving Technologies for Blockchains" - `http://www.r3cev.com/s/R3_Confidentiality_and_Privacy_Report.pdf`

[33]A technique that allows a party to generate many "single-use" public-key addresses, thereby hiding the identity of the recipient of a transaction.

[34]A method of creating transactions where the sender cannot be identified by their public-key.

like Pedersen commitments[35] or the enigma protocol[36], could hide transaction or computation data. Implementation of these technologies in blockchains may be recent, but the underlying cryptography is mature and well understood. More recent cryptographic techniques like zero-knowledge proofs[37], implemented in the fully-anonymous cryptocurrency Zcash, allow for greater privacy and confidentiality guarantees.

Privacy and confidentiality-protecting techniques such as anonymizing data on a distributed ledger is an active area of research. It is possible that some combination of these techniques could offer sufficient protections that, for instance, data stored on a blockchain would no longer be considered "identifiable" within the meaning of regulations like the GDPR.

For now, the comparatively low-tech solution of limiting read-access may be the only practical method for financial institutions to meet their legal obligations on data treatment.

Further, in some cases even mature cryptographic techniques may not be a sufficient solution. While some cryptographic methods may be essentially impossible to break with today's technology, that may not be true in 10, 20, or 50 years. The recent discovery of a collision in the SHA-1 hashing algorithm, effectively rendering it insecure and obsolete, served as a useful illustration of this general problem. Over time, other hashing algorithms will be compromised as well, as they have been in the past.[38]

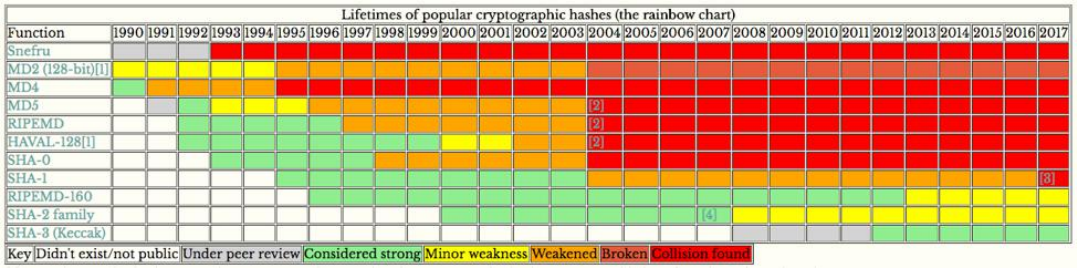


Figure 5: Lifetimes of popular cryptographic hashes (`http://valerieaurora.org/hash.html`).

Data that is simply released into the public may eventually become insecure within a timeframe that, while it may seem remote to the average person, poses legal or other risks to large financial institutions and their customers. The rise of quantum computing poses a related risk, as it may render some cryptographic techniques less useful in the coming decades.[39]

# 3 Reporting of OTC Interest Rate Swap Transaction Data Using DLT

In this section, we examine in detail how transaction reporting for an OTC Interest Rate Swap could function on the Corda platform. We will focus in particular on the reporting requirements under US law, though many of the processes described could be adapted to similar requirements in other jurisdictions.

## Introduction to Corda

Corda is R3's permissioned distributed ledger platform designed for use in financial services.[40] Corda is not a blockchain - there is no series of hash-linked blocks. Rather, Corda adopts some of

[35] A technique that can be used to disguise the amount of cryptocurrency being transferred.

[36] A proposed method of enabling distributed secure multiparty computation, described in a whitepaper published in June 2015. In theory, Enigma would allow a blockchain like Ethereum to process scripts (i.e. "smart contracts) without revealing the data being processed - see: G Zyskind, O Nathan, A Pentland.

[37] More specifically, "zero-knowledge succinct non-interactive arguments of knowledge" or zk-SNARKs, which are used in Zcash. Zero-knowledge proofs themselves have been studied since the 1980's.

[38] For clarity, cryptographic hashes are not primarily used to protect privacy in blockchain systems – rather, they are used to ensure (among other things) integrity of data – i.e. linking blocks by their unique hashes makes it very easy to verify that data in previous blocks has not been altered. The compromise of the SHA-1 algorithm is an example of a general problem that has privacy implications – that what is secure today may not be secure tomorrow, and we expect data stored on a public blockchain to be around for a very long time.

[39] For a discussion of this point, see: D Yang, J Gavigan, Z Wilcox-O'Hearn (2017), Appendix 5.1.

[40] For more information, see the Corda Technical Whitepaper: `https://docs.corda.net/_static/corda-technical-whitepaper.pdf`

the design choices of blockchain technology while dropping others.

In Corda, there is no global state shared by an entire network, like in Bitcoin (a ledger of all bitcoin) or Ethereum (the state of an abstracted virtual computer). Rather, information is only available to certain entities - like counterparties, regulators, and other nodes where they are necessary.

This means that "consensus" is achieved only at the level of a specific financial agreement, rather than across the whole network. Consensus on the validity of any given transaction - i.e that it follows defined rules and contains the proper cryptographic signatures - is accomplished by having each party to the agreement validate the transaction. Consensus on the uniqueness of any given transaction - i.e. that it is not spending funds that have already been spent - requires a third-party observer. Special nodes called "notaries" play this role. A notary could be a network of nodes running a consensus algorithm (i.e., similar to existing public blockchains) or a single known entity given legal responsibility for performing the notary function.

Financial agreements are represented by data structures called "state objects" which contain both legal prose and code. In other words, Corda facilitates "smart legal contracts" that define a legal relationship between parties and take advantage of smart-contract code by automating processes related to that financial relationship.

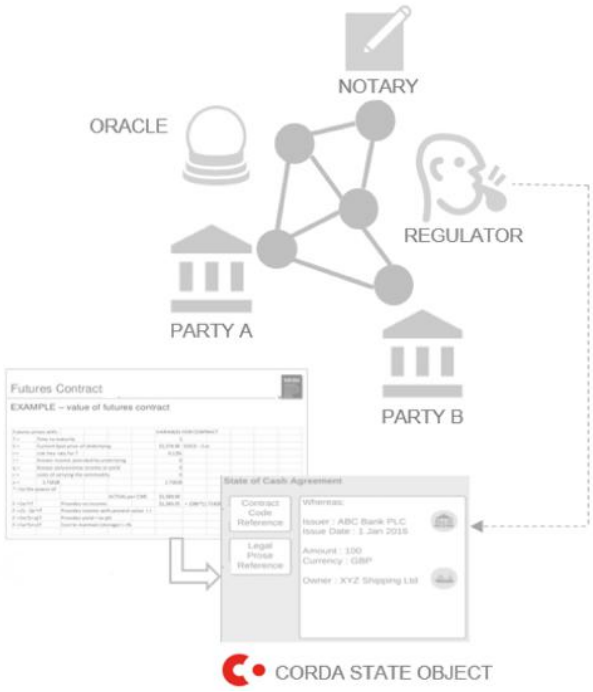


Figure 6: A state object representing a financial contract and group of participants, including parties, regulator, oracle, and notary node.

As described above, only certain parties will have access to a state-object representing a financial agreement. In particular, Corda is designed to allow regulators to operate[41] nodes that are given access to the state of financial agreements over which that regulator has a supervisory role and, therefore, a need to see certain information about certain transactions on the ledger. Whether information is sent to a regulator node is determined by the "flow framework" of a specific Corda application (or "CorDapp"). These flows are designed by the authors of individual CorDapps, allowing flexibility in how information is shared for different types of financial relationships. If desired by the parties, a flow framework will exist that pushes necessary data to regulator nodes when appropriate. These "regulator nodes" are the basis for the reporting process we describe below.

[41] Or have nodes operated on their behalf by a third party.

## Regulatory Reporting for OTC Swaps

Interest-rate swaps (IRSs) are contracts whereby two parties exchange streams of interest payments. In a "vanilla" IRS, one party makes interest payments on a fixed rate, while the other makes payments on a floating rate (fixed to some standard rate, e.g. LIBOR). Over a defined period each party makes payments to the other according to a set schedule.

Post-financial crisis regulatory reforms like title VII of Dodd-Frank introduced new reporting obligations for OTC derivatives. Previous to these regulations, OTC derivatives were typically traded between individual firms and were not required to be reported to a regulator.

The reporting requirements we consider here are those for OTC interest-rate swaps, which are governed by the Commodity Futures Trading Commission (CFTC). There are two components to these regulations.

- First, the real-time public reporting of swap transaction data under 17 CFT § 43[42]. This data must be reported to a "Swap Data Repository" (SDR), a designated entity who anonymizes the data and releases it publicly. This information is reported in close to real time, as transactions are executed.
- Second, the more general SDR reporting requirements under 17 CFR § 45[43]. This data is reported to a regulator. Generally, this means that the data is again sent to an SDR, where it is made available to regulators.

In our analysis below, we will consider how the reporting requirements under these two sections could be satisfied through an IRS transaction on Corda.

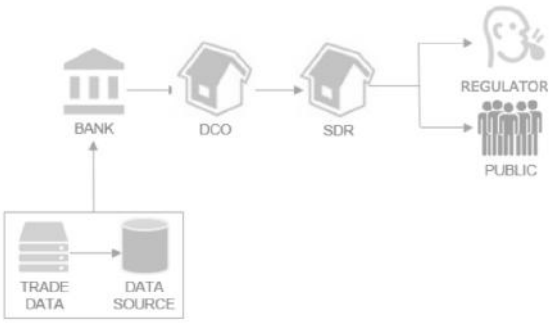


Figure 7 - A generic overview of the current process of reporting transaction data on centrally cleared OTC trades.

## Regulatory Reporting Workflow in Corda

### Pre-trade Transparency and Execution

We assume that the parties already have an ISDA Master Agreement and Credit Support Annex (CSA) in place, setting out the basic terms and collateral arrangements between the parties for all transactions of a certain type, including interest-rate swaps. Further, the example described below focuses on bilateral transactions – centrally cleared transactions are discussed on page 27-28.

Our two parties – A and B – wish to create an IRS. Both parties have access to an implementation of the Corda platform, and each maintains a node (or more likely, several nodes) that allow them to form and enter into agreements.

The parties would use template agreements[44] for their IRS and negotiate through a user-interface connected to the Corda platform. The offering and final acceptance of terms would be managed through a Corda "flow framework"[45] - a protocol for communication between parties.

[42] Real Time Public Reporting, 10 CFR § 43 2012.

[43] Swap Data Recordkeeping and Reporting Requirements, 10 CFR § 45 2012

[44] See for instance C D Clack, V A Bakshi, and L Braine (2016). "Smart Contract Templates: foundations, design landscape and research directions" - `http://www0.cs.ucl.ac.uk/staff/C.Clack/SCT2016.pdf`

[45] `https://docs.corda.net/key-concepts-flow-framework.html`

As part of this flow, a Corda transaction is formed that captures the terms of the IRS agreement between the parties. The transaction is signed and validated by each party's node, and a notary is designated that will be responsible for maintaining consensus over future updates.[46]

In this case, there is also a regulator node on the network. As defined by our flow, this node is included in the set of nodes who are able to see details of the IRS between A and B. The initial transaction is broadcast to the regulator node, as well as the confirmation from a notary that the transaction has been approved. Our flow also specifies that any updates to the state of the IRS at a later time will also be available to the regulator node. For the purposes of this analysis, we assume that the regulator node is operated by a SDR.

Once the transaction has been validated and signed by each party's node and our flow is completed, a state object is created representing the IRS, which includes references to the ISDA Master Agreement and CSA. At this point, the parties have entered into the IRS, and reporting obligations are triggered.[47]

### Post-trade Transparency

**Swap Transaction Reporting Requirements - "Section 43 Obligations."** CFTC regulations require certain information to be reported in real-time to a SDR.[48] For instance, this includes the time and date of the execution, an indication of whether the swap is collateralized, start and end dates, settlement currency, asset class, and other details about the transaction.

Because our regulator node, operated by an SDR, has had access to this transaction from its beginning, the majority of these reporting requirements will have already been met. The flow framework designed for our IRS specifies that when the IRS is created, the necessary data is "pushed" to the SDR node. Both the parties and the regulator node will have received the notary-signed transaction as a matter of course, and used it to update their copy of the shared state. Rather than receive a separate report attesting to the facts of the transaction between A and B, the SDR has an authoritative copy of the contract itself.

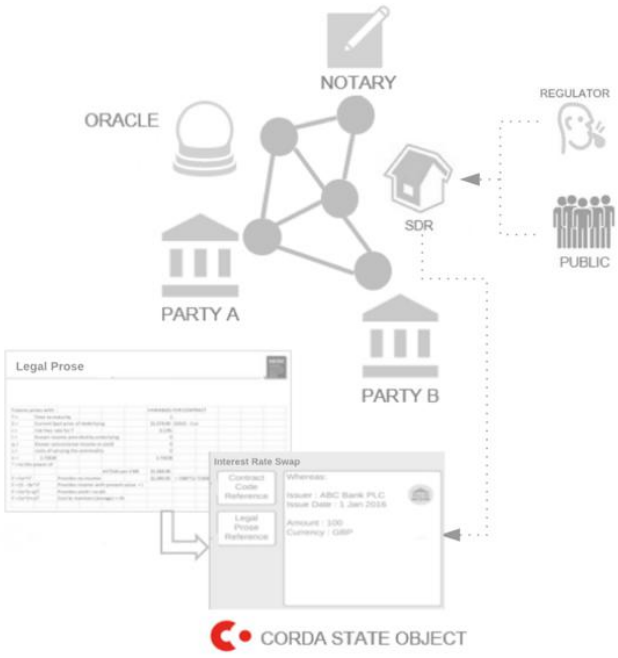


Figure 8 - Regulator node operated by SDR has access to state object representing transaction, fulfilling reporting obligations.

[46] In Corda, notary nodes are used to maintain consensus. A notary checks that any update to a state-object is unique - in other words, that it consumes only state-objects that have not yet been consumed. The notary is not required to ensure uniqueness in this transaction because we are "issuing" a new asset - there is no prior state to consume.

[47] We assume that the act of cryptographically signing the transaction constitutes "execution" in the meaning of the Act.

[48] For the full list of required information, see Appendix A to Part 43: `https://www.law.cornell.edu/cfr/text/17/appendix-A_to_part_43`

However, some information required under Part 43 may not necessarily be included in a standard IRS agreement. For instance, most swaps are subject to a mandatory clearing requirement, meaning that they must be cleared by a Derivatives Clearing Organization (DCO). However, there are exceptions to that requirement. Part 43 requires that any trade not centrally cleared must indicate in their report the legal exemption they are relying on. More generally, it may simply be the case for many reporting obligations that there will be data that must be reported that would not, ordinarily, be included in the legal contract or confirmation letter for the relevant transaction.

Fortunately the solution is simple: in order to ensure that the relied-upon exemption is reported to an SDR, we would simply include this and any other information in the contract itself, or as an attachment to the transaction[49] This means that the SDR will have received this information along with all of the terms of the contract.

To complete the Part 43 reporting process, the SDR would extract the necessary information from the state-object, anonymize it, and report it publicly through the same process used today.[50]

One component of the real-time reporting obligations under Part 43 offers an interesting opportunity for automation and, perhaps, more finely-tuned regulations to promote more efficient markets. Part 43 includes a rule whereby certain "block trades" of sufficiently large value are granted a time delay in when they are reported publicly.[51] The rationale for this rule is that real-time reporting of these large trades can actually reduce market liquidity and disproportionately impact the market.[52]

Imagine that Corporation A raises capital by issuing a fixed-rate bond. The corporation may choose to hedge its interest rate risk by entering into an IRS with Bank B. If the fact of this large IRS is publicly reported, it is a signal to the market that someone (Bank B) will be seeking a way to hedge its risk from that transaction. This can lead other entities to adjust pricing in anticipation of a trade, increasing costs to Bank B. This may lead Bank B to not enter into the trade in the first place, or to raise their own prices to Corporation A, both of which might inhibit investment in the economy.[53]

A reporting structure like the one described above offers opportunities to automate this process. The regulator node would automatically check the value of any given IRS, and according to predefined smart-contract code, trigger a reporting-delay if the value exceeds the current minimum block size. Further, we can imagine that this could enable regulators to fine-tune reporting delays according to transaction size, as opposed to a blanket "minimum" triggering a delay. Assuming there is some function that expresses the relationship between trade size, reporting delay, and impact on market liquidity, then we could build that function into our rule, triggering the appropriate reporting delay automatically for trades of every size.

**Swap Data Record Keeping and Reporting Requirements.** Under 17 CFR § 45, parties must report information in two categories: (1) Creation data and (2) Continuation data.

**Creation data** includes the Primary Economic Terms (PET) and confirmation data. The PET of an IRS include information like asset class, execution venue of the swap, start and end dates, price, and other information.[54] Confirmation data includes all of the terms of a swap (strike price, notional amount, contract type, etc.) agreed upon by the parties, as well as certain identifying information for cleared swaps. Like under Part 43, this information must be reported to an SDR as soon as technologically possible. As with our description above, this information will have been automatically pushed to an SDR by virtue of the SDR node's inclusion in the flow framework for this transaction.

**Continuation data** is reported throughout the life of the swap, in order to make sure that the information held by the SDR stays current and accurate. This data has two components. First, life cycle[55]data. This includes any event that alters the PET. For instance, if the terms of the swap are amended by the parties. Because the SDR node is part of the flow framework for this transaction, they will again automatically be aware of any alteration to the PET. Second, valuation data. Per

[49]Corda "attachments" are discussed on p. 17 of the technical whitepaper: `https://docs.corda.net/_static/corda-technical-whitepaper.pdf`

[50]Real Time Public Reporting, 10 CFR § 43.4(d) 2012.

[51]Real Time Public Reporting, 10 CFR § 43.6 2012.

[52]"*Block trade reporting for over-the-counter markets." International Swaps and Derivatives Association*, January 18, 2011 - `http://www.isda.org/speeches/pdf/block-trade-reporting.pdf`

[53]This example adapted from "Block trade reporting for over-the-counter markets." p6.

[54]Swap Data Recordkeeping and Reporting Requirements, 10 CFR § 45 Exhibit C 2012.

[55]Or "state data" for snapshot method. We assume life cycle data for the purposes of this section.

the regulations, this is "all of the data elements necessary to fully describe the daily mark of the transaction."[56] This data must generally be reported to the SDR daily.

Reporting daily valuation data is more complex than life cycle data. Reporting changes to the IRS is simple because the SDR already has access to the state, since our flow says that any update is sent to the regulator node. But valuation data requires us to design a separate process that can incorporate outside information (for instance, to draw market data from an oracle). We should also consider that each party to the IRS may not need to receive daily valuation data, only the regulator node.

One way to achieve this would be to create a separate "valuation state object". This state object would refer to the IRS state object, and be incorporated into the flow framework we used to create and manage it, but is a separate data structure. Like the IRS state object itself, the valuation state object would be shared with the regulator node. By using a separate state object, we avoid the problem of having to "update" the IRS state object every day with new valuation data.

Each day, the reporting party tasked with submitting valuation data would form a new Corda transaction.[57] This transaction would take the existing valuation state-object as an input, and output a new valuation state-object. The transaction would essentially update the valuation seen by the regulator node each day, with the information necessary to disclose the daily mark of the swap, which will include the methodology and assumptions used to prepare the mark.[58]

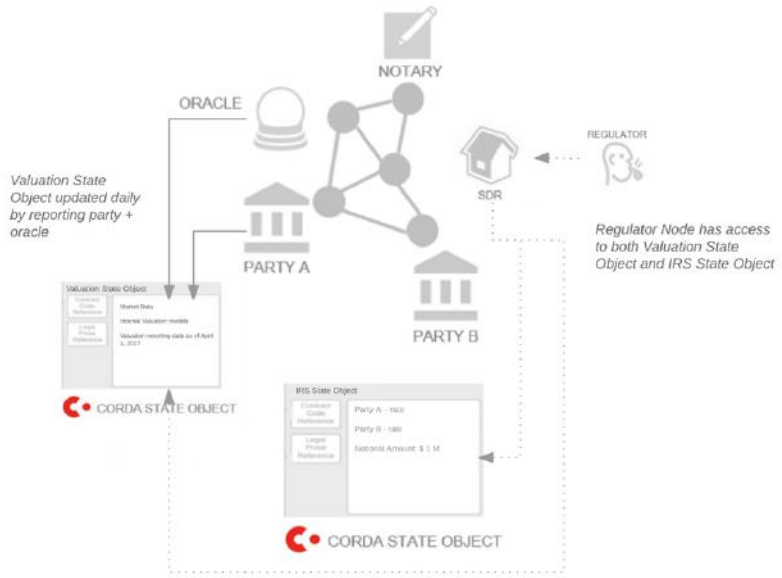


Figure 9 - Regulator node operated by SDR has full view of Valuation State Object, updated daily to satisfy § 45 valuation data reporting obligations.

## Additional Considerations

### Centrally Cleared Swaps

Under CFTC regulations, certain types of swaps must be "centrally cleared".[59] This means that the ultimate counterparty to both sides of the trade is a specialized institution - a Derivatives Clearing Organization (DCO). In other words, the DCO is the buyer to every seller and the seller to every buyer for certain types of derivative.

For centrally cleared swaps, the process outlined above would be slightly different. The terms of the IRS negotiated between the parties would be the same, but the final agreement represented by the state-object would be broken into two transactions, each facing a third party, the DCO. Like our other parties, the DCO would operate a node on the Corda platform for this purpose.[60]

[56] Swap Data Recordkeeping and Reporting Requirements, 10 CFR § 45.1 2012.

[57] On Corda, and more generally in DLT systems, a "transaction" is a proposed update to the ledger. In this example, our transaction is an update to our valuation state-object.

[58] Defined as the "mid-market mark" by Disclosures of Material Information, 10 CFR § 23.431 2012.

[59] Clearing Requirement and Related Rules, 10 CFR § 50.4 2012.

[60] This topic is the subject of an upcoming paper written in collaboration with R3: "Implementing derivatives

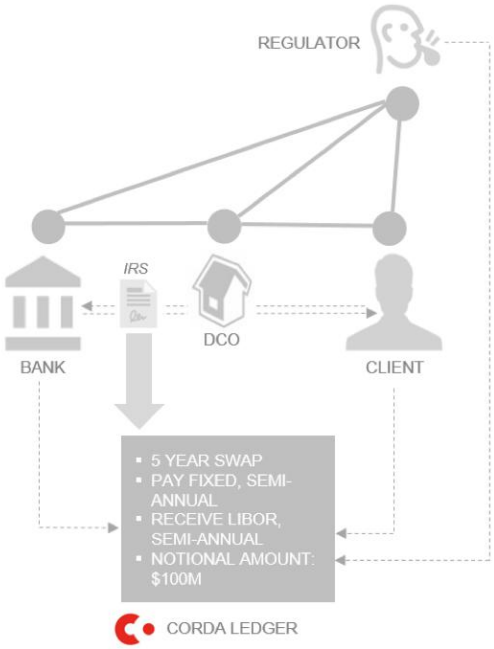


Figure 10 - A centrally cleared swap, where a DCO sits "between" two parties to an IRS.

**Reporting Party Responsibility**

Parts 43 and 45 set out rules determining who is responsible for reporting specific information. Under Part 45, any swap that is subject to mandatory clearing through a DCO will in turn have its reporting obligations rest on the DCO itself.[61] This includes the majority of IRSs, meaning that in practice the parties to any given IRS will likely not be responsible for reporting, as the DCO will be the reporting party.

For swaps that are not centrally cleared, one or the other of the counterparties to the IRS will be responsible for meeting the reporting obligations laid out above[62] In practice, if these IRS were traded over the Corda platform, this would not significantly change the process described above. Updates to the state object (creation data and life-cycle data) will be seen by all parties included in the flow framework and, from the perspective of the reporting party, happens automatically without requiring any special effort on their part (beyond agreeing to the update itself). In practice, it is the defined flow framework between the parties that ensures information is available to the regulator node. Nonetheless, the reporting party would be held responsible if the reporting process somehow fails.

The exception to this is valuation data, as described above. Here, the reporting party (in the case of non-cleared swaps) must carry out an additional process (though it may be automated) by signing transactions to update the IRS valuation state object described above.

Under Part 43 if a swap is executed on a registered Swap Execution Facility (SEF) or a Designated Contract Market (DCM), then the parties to that swap will have met their reporting obligations.[63] Otherwise, the swap would be considered an "Off-facility swap", in which case reporting obligations are determined by the rules described in Part 43.3(a)(3), unless otherwise agreed to by the parties. If we assume that a trade over Corda is an off-facility swap, then one or the other of the counterparties will be responsible for fulfilling the real-time public reporting obligations. Again, as described above, transaction data reporting will happen automatically from the perspective of the reporting party, as information about each swap is shared with a regulator node per the flow framework.

---

clearing on Distributed Ledger Technology Platforms", by Colin Platt, Massimo Morini, and Péter Csóka.

[61] Determination of which counterparty must report, 10 CFR § 45.8(i) 2012

[62] Determination of which counterparty must report, 10 CFR § 45.8 2012. However, this does not include daily valuation data where the reporting counterparty is neither swap dealer nor a major-swap participant. In that case, the reporting counterparty must report the current daily mark each fiscal quarter (Swap data reporting: continuation data, 10 CFR 45.4(d)(2) 2012).

[63] Real Time Public Reporting, 10 CFR § 43.3(a)(2) 2012.

# Conclusion

Distributed ledgers are a suitable technology for achieving many of the basic processes required for regulatory compliance. Allowing multiple parties to access a shared authoritative record offers an elegant simplification of many regulatory functions. Removing the need to send copies of information may result in fewer errors, improving the quality of data used by financial institutions to model risk. The built-in features of distributed ledgers - like requiring transaction validation at each step, and recording a full audit trail complete with cryptographic signatures - improve data integrity throughout the entire process.

Applying this process to transaction reporting of OTC interest-rate swaps, we saw in practice where a distributed ledger could offer advantages. Most information that must be reported under CFTC regulations can be easily shared with a regulator. Because we are sharing an authoritative record of the IRS itself, most data (for instance, the Primary Economic Terms) will be reported automatically, simply by giving a regulator node access to the the state-object through the flow framework.

In other cases, parties may need to design additional processes to meet their reporting obligations. Information required by a regulator might not, in the normal course, be included in the text of a contract. For instance, under Part 45 a party must report whether they are relying on a particular exemption to avoid the mandatory clearing requirement. To be reported, this information must be included in the state object representing a financial agreement.

Some rules could be improved by being embedded directly into the platform as smart-contract code. For instance, the reporting delays for block trades of a certain value could be triggered automatically when reported to an SDR.

The use of distributed ledgers could, in theory, reduce the need to rely on entities like SDRs who serve as middlemen in the regulatory reporting process. Advanced privacy-preserving techniques like those described in Part 1 could be used to anonymize data before it is released to the public, rather than rely on the SDR to carry out this function. Regulators may no longer require the SDRs to serve as a central store of transaction data, if all transaction data is stored on a distributed ledger to which they have access.

# References


[1] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten. Research perspectives and challenges for Bitcoin and cryptocurrencies. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2015.

[2] J. Brekke and E. Haase. Satoshi Oath, 2016. `http://ipfs.b9lab.com:8080/ipfs/QmXysWEAexXQqYZhTGpECvksnaBkSEWHdGhM7vNeHxue2g/`.

[3] M. Chase and S. Meiklejohn. Transparency overlays and applications. In *Proceedings of ACM CCS 2016*, 2016.

[4] G. Danezis and S. Meiklejohn. Centrally banked cryptocurrencies. In *Proceedings of NDSS 2016*, 2016.

[5] I. Eyal, A. E. Gencer, E. G. Sirer, and R. van Renesse. Bitcoin-NG: a scalable blockchain protocol. In *Proceedings of NSDI 2016*, 2016.

[6] A. Gervais, G. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun. On the security and performance of proof of work blockchains. In *Proceedings of ACM CCS 2016*, 2016.

[7] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2016.

[8] V. Lehdonvirta. The blockchain paradox: Why distributed ledger technologies may do little to transform the economy, 2016. `http://blogs.oii.ox.ac.uk/policy/the-blockchain-paradox-why-distributed-ledger-technologies-may-do-little-to-transform-the-economy/`.

[9] L. Luu, V. Narayanan, K. Baweja, C. Zheng, S. Gilbert, and P. Saxena. A secure sharding protocol for open blockchains. In *Proceedings of CCS 2016*, 2016.

[10] A. Miller, A. Juels, E. Shi, B. Parno, and J. Katz. Permacoin: Repurposing Bitcoin work for data preservation. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2014.

[11] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi. Town Crier: an authenticated data feed for smart contracts. In *Proceedings of CCS 2016*, 2016.