CAPCO

COMPLIANCE AT SCALE

OVERCOMING LEGAL BARRIERS TO DATA DEMOCRATIZATION

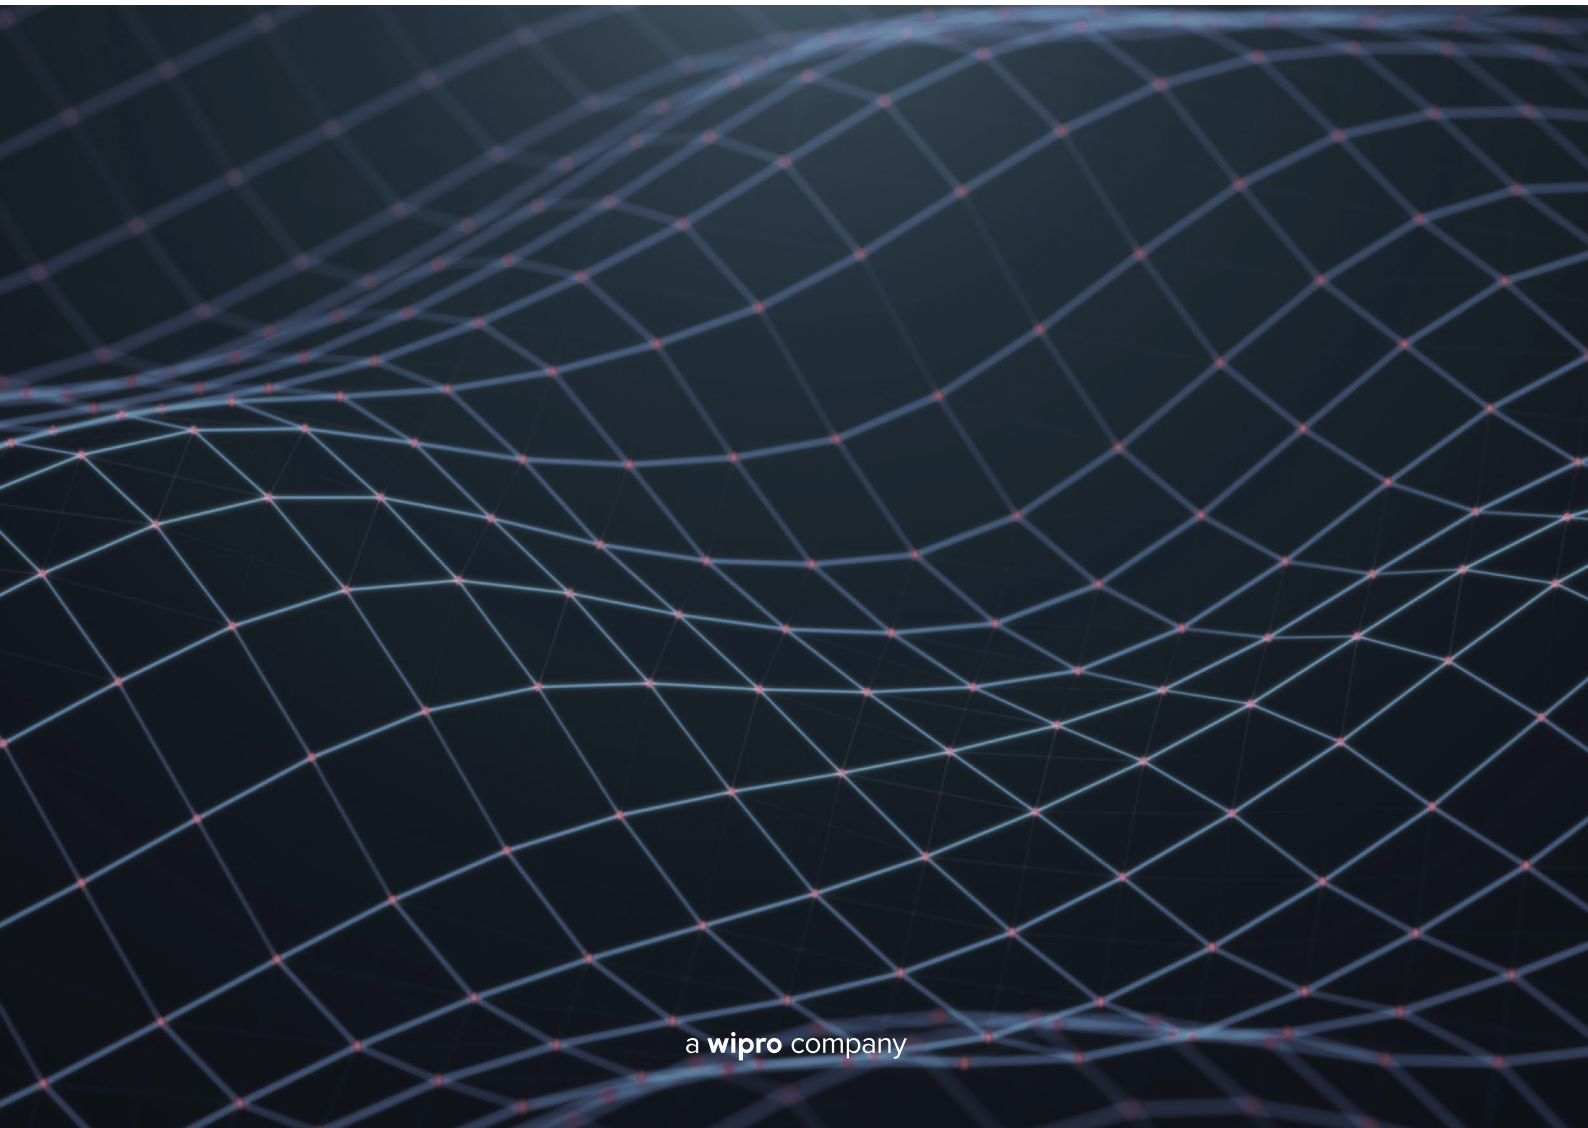

a **wipro** company

COMPLIANCE AT SCALE
OVERCOMING LEGAL BARRIERS TO DATA DEMOCRATIZATION

As organizations advance on their digitization journeys, their growth strategies, decision making and operational optimization are becoming more data-driven – and that data is becoming more centralized and co-located. However, this conflicts with a regulatory trend (especially in relation to data privacy and protection) that requires the minimization of data access and sharing.

Given the inherent complexities of large, multinational and multi-line enterprises, such organizations face common challenges in managing and accelerating access to data to meet the business' growing data demands while also remaining compliant with approved data practices.

This paper highlights those challenges and looks at how organizations might achieve effective compliance while also avoiding the significant scaling-up of resources or the creation of burdensome control processes that typically result from data privacy and protection implementations.

In today's digital world, organizations are increasingly seeking ways to leverage value from their data assets to optimize operational costs and broaden revenue streams. Central to this is the increasing application of data science and advanced analytics to gain insights into the digital mechanisms that drive the economy as well as their own organizations.

To support this, many financial services organizations (FSOs) are actively improving the data literacy of their employees while simultaneously migrating to low cost, scalable and more accessible data storage solutions, including cloud-based platforms. This trend of data 'democratization' is allowing these organizations to shift gradually from having pockets or silos of advanced data usage to enterprises characterised by global data access.

The drive for data democratization also highlights a commensurate challenge – how do financial services organizations (FSOs) ensure they adhere to local, regional and global regulations as well as internal policies (not just in the data privacy and protection space) in a fashion that is at once effective at scale and in alignment with a rapidly changing legislative landscape?

The challenge of compliance at scale

Enterprise data warehousing and the push towards data democratization reflect a common and long-term strategic trend towards data co-location – the hosting of data originating from multiple locations in a single, or reduced number of, physical or logical locations. Conversely, the core principles of data privacy and protection require data to be more ringfenced and less freely accessible. The challenges arising from this dichotomy are becoming more significant for FSOs.

Some of the key problems they face include:

- Adapting to a complex and changing legislative landscape while avoiding a one-size fits all approach
- Understanding the legislation applicable to large data sets aggregated in cloud environments
- Reducing the cost of compliance while simultaneously remaining effective in mitigating risk
- Managing the increasing control burden in ways that allow innovation to thrive.

When combined with factors such as the efficiency of control design, effectiveness of embedment, and resistance to change, it is easy to see why data-related regulatory fines are on the rise as cloud and analytics platform implementations continue to flourish. H1 2022 has already seen $100m in fines issued, representing a 92% year-on-year increase.[1]

Against the backdrop of increasing size and severity of fines issued, an eye-watering amount is still being spent on data privacy and protection compliance. According to one study, a total of $1.2bn has been spent on compliance alone in the UK – yet with a questionable impact on effectiveness, as nearly 75% of UK companies still do not comply with GDPR data request requirements.[2] That so much is still being spent on compliance with little improvement relative to the investment underlines the need to address the challenge in ways that are significantly more cost-effective.

Addressing the challenge

If FSOs are to become data compliant at scale, it is not merely a case of achieving regulatory compliance, but doing so in a manner that is designed around economies of scale. This means that an organization's control environment, aligned to its data management capability and strategic data architecture, drives embedment of compliance controls with minimal additional overheads. That is, they are already incorporated in the management procedures and architectural principles.

Existing investment in data management capabilities represents an opportunity to address regulatory compliance to data sharing and processing in a manner that is both lean and builds on the foundations that should already be well established. The right operational solutions and tooling can enable rapid, scaled adoption, meeting the challenge of providing compliant access despite the accelerating demand for data and a regulatory environment that will continue to change at pace.

Furthermore, compliance at scale depends on how the control environment is intrinsically linked. Scale returns are driven by the realization that data privacy and protection controls such as the Data Protection Impact Assessment (DPIA) and the Record of Processing (ROP) do not exist in isolation. They are dependent both on each other, and on the way in which an organization manages its data.

Even for a relatively simple business model, the challenges of addressing data sharing and processing compliance can be daunting. However, a successful business will grow and inevitably must manage the changing and increasing complexity this growth would bring. We will set out approaches across data management and the control environment that leverage existing foundations to help prepare a business for future growth and the data protection risks, as well as the opportunities that subsequently arise.

BUILDING THE DATA MANAGEMENT FOUNDATIONS

Strong data management is the foundation of an organization's ability to understand what data it has in its possession, where that data is stored, how it flows through its IT systems, and for what purposes it is being processed. Driven by regulatory directives like BCBS 239, CCAR and Sarbanes Oxley, or the commercial advantages in maintaining data-driven businesses, many organizations have capabilities in place to maintain the integrity of, and build trust in, the data they hold.

These requirements of a Data Management capability that enable it to be effective also serve a dual purpose in empowering privacy professionals within an organization to better manage the data assets: understand what sensitive data they have and where; as well as managing the risks associated with them.

Historically, data management evolved from its origins as a mere reaction to legislative requirements to being the cornerstone for which an organization establishes trust in its own data – ensuring that the right data of the right quality is being used for a specific purpose. In this respect it has gone from a tactical consideration to a truly strategic capability. Similarly, the approach for maintaining compliance with data privacy and protection legislation must shift away from standalone tactical projects to a fully embedded process that works in concert with existing strategic data management framework and procedures.

When data privacy remediation projects are addressed through a siloed, one-off approach, this assumes that this relatively new legislation is mature and will neither change nor the impacts of which be better understood as demands on data are further enhanced. To avoid this, the procedures to maintain compliance should be designed alongside existing controls that already have flexible designs to accommodate differing demands and changes in the regulations. As well as ensuring that privacy risk is managed throughout all data activities, this also delivers greater efficiency through repeatable processes.

Setting targets

Organizations should strive to meet key principles for data collection, sharing and usage controls. In summary, these are:

- Lawful basis – the first principle of data privacy law that ensures that all personal data is processed lawfully, fairly and transparently
- Ethical basis – asks whether how data is being shared and processed is the right thing to do or whether damage could be done to the data subject directly or indirectly
- Minimization and limitation of data storage and processing – identify and use only the relevant and least amount of information required for the required processes
- Third-party risk management – what data do third parties have access to, do they have the necessary security and privacy protocols in place and what would the impacts be if there was a breach?
- Accuracy, integrity and confidentiality – is the quality of the data high and updated to ensure it is accurate? Have technological and organizational measures been taking to ensure the security of data
- Governance and accountability – does the organization have in place technical and organizational measures in place to implement effective data protection controls; is it clear who takes ownership for the data and the controls and protections in place?

There are a number of data governance assets that directly enable these principles:

Data catalogue – a foundational asset that captures data definitions, classification (public, internal, sensitive, MNPI, PII, etc), purpose of use, risk impacts and ownership at an elemental level:

- Allows privacy professionals to understand the risk associated with the data in an organization
- Working with data governance, access entitlements can therefore be easily controlled.

Metadata tagging – including a category related to the purpose of processing of the data:

- Allows the privacy professional to understand how all data is used
- Can be leveraged to facilitate automation
- Provides an enhanced understanding of potential regulatory conflicts in the sharing and movement of data.

Data Lineage flows – data managers want to understand the sources of their data, where that data is being moved to and from, and how it is potentially transformed through to the end consumers of that data:

- The privacy professional within the organization needs a complete picture of ownership and authority of the data elements through to where it is then shared and processed.

In addition to these assets, a mature data management capability should have in place existing processes, roles, responsibilities, and accountabilities which can then be extended to data privacy, collection, sharing and usage controls. This can be applied through a centralized, global governance structure that is already defined and mobilized.

Compliance to the principles for data collection and usage are then better able to be achieved:

- Lawful basis can be determined using metadata tagging and classification of the data
- Ethical basis for the processing of data can be governed through embedded controls and procedures
- A well maintained data catalogue and governance process can ensure that only the necessary amount of correct data is shared and processed
- The data lineage flows will also inform where data is moving to third parties or other entities to allow the organization to determine the potential risks
- Data quality metrics related to critical data will also be understood through the data catalogue and strong controls and monitoring, further ensuring compliance to regulated data management and privacy principles.

The realities of data management for compliance at scale

An optimized and established data environment into which data privacy compliance controls can be incorporated is a rare thing indeed. Consequently, it is important organizations consider approaches for accelerating their data management capability.

The enterprise-wide operational risks for organizations – and the management of those risks through controls and monitoring

– should be the basis for which data protection-specific controls are either expanded upon or introduced. This enables implementation into a highly mature and embedded operational framework, thereby minimizing additional resourcing overheads or significant procedure and control monitoring rework.

As with the integration of privacy control processes with those of data management, change management frameworks, tollgates, and metrics should reflect and integrate alongside other data and metadata within an organization's business lines and functions (sales and marketing, HR, information security, etc). Together with the data management capabilities' view of data quality and lineage, this allows for greater insights into where the risks lie, which business processes are exposed, and the root cause of those exposures, enabling a risk view that allows for faster determination of what risks can be accepted.

A strategic, end-to-end approach for data privacy compliance at scale would require the following three-step approach:

1. Discover / identify data – what data does the organization have across its estate, and where?

2. Classify data – capture metadata on the data (preferably at point of ingestion); this allows for effective classification of the data, enabling improved understanding of the data asset and the applications of automated governance procedures

3. Control – leverage capabilities such as knowledge graph technologies to apply controls based off the tagging and classification of the data being utilized.

Intelligent discovery

The data catalogue is one of the most powerful assets in respect of data management and data privacy risk management. However, the completeness and quality of that catalogue can vary. The risk is that, even with well-designed controls, procedures and effective implementation within a mature data management operating model, the catalogue can be undermined by gaps in the assets that enable effective assessment of risk. To address this potential gap or validate current data management information, an intelligent discovery and scanning tool can be leveraged to accelerate the cataloguing of data, identifying sensitive data and thereby enabling privacy management.

Such a tool should be used to identify the locations of sensitive data across the organization's ecosystem. In doing so, it can address the probability of risk in the areas of the business it is deployed, together with assessment / validation of the data lineage, control mapping and cataloguing already captured. This should also include scanning for unstructured data across multiple different locations and (once its distribution is understood) which controls around sharing need to be applied. This would include how sensitive data is shared across emails and other forms of communication; where it is located in simple forms or documents that exist in file directories.

This exercise (alongside ongoing, periodic scans) delivers analysis of the areas of risk across the organization's environment and also enables a business to execute against its obligations for Data Subject Requests (DSR) requests, or right to be forgotten.

When considering the need to manage sensitive data across multiple domains in a large-scale organization, manual review and maintenance simply becomes untenable and an inherent source of risk. Ongoing intelligent scanning provides a means to ensure that regardless of the velocity of digital change within an organization, it is possible to remain compliant as well as identify potential risks before they directly impact the business and require remediation.

Metadata capture and classification of data

The quantity and complexity of global data regulation necessitates an increased level of control automation in order to facilitate decision-making around data sharing / access compliance while also reducing the need for human involvement. This requires structured inputs (metadata) in order to drive a rules-based, machine-led approach, minimizing the need for human input – see next section on Scaling the Control Environment. Strategically the metadata should be consumed from a single authoritative metadata source to minimize the cost and complexity of metadata management (e.g. the Enterprise Data Catalogue).

Basic metadata concepts that should be captured to facilitate regulatory compliance include:

1. Data Class – the type of data defined in the Enterprise Data Model (e.g. retail customer data)

2. Regulatory Data Category Flag – a flag applied to a data set / record / application to indicate whether it contains a specific type of data of regulatory concern in relation to the collection, storage, transfer, access or processing of that data (e.g. personal data)

3. Policy Flag – a flag applied to a data set / record / application to indicate the data regulations it is subject to (e.g. privacy)

4. Controlling Legal Entity – the organization's legal entity that has legislative or regulatory authority over the collection, storage, transfer, access or processing of a given set of data (e.g. UK retail bank).

Note: these are not definitive (especially in the context of broader data governance and the data catalogue asset). The expectation is that additional metadata would be captured already (e.g. the application that data is stored on, purpose of processing).

Ideally metadata is captured at the point of ingestion. However, in reality most organizations will not have done this and have a legacy of data that needs to be retrospectively augmented. Organizations that have taken a manual approach to tagging and classifying the data find this to be resource intensive and a difficult operational activity to maintain with high levels of quality.

Metadata harvesting can be accelerated using similar intelligent scanning solutions as those used for data discovery.

It should be emphasised that the use of an intelligent scanning tool for the purpose of initial discovery and maintenance of sensitive data and its metadata would best deliver to its objectives through a service capability. This is to provide structure and governance to the process – ensuring not just quality of the outputs, but also comprehensive and prioritised coverage across the enterprise.

SCALING THE CONTROL ENVIRONMENT

The control at scale challenge

Different jurisdictions have different data protection standards and sensitivities which has left organizations caught between the rock of enforcing common global standards and the hard place of achieving compliance in a manner tailored to specific legislation in a specific jurisdiction.

In the face of such complexity, the challenge increasingly becomes delivering the simpler – but often manually intensive – goal of basic regulatory compliance in a way that can be tailored to local laws and sensitivities, but without disrupting the business and stifling innovation. Finding the correct balance between operational effectiveness and compliance will deliver a significant advantage over competitors. Achieving this balance requires recognition from both sides, as compliance for a use case is meaningless outside of a business driver. Both sides of this symbiosis must look to accommodate the other in order to find the correct balance of constraints on the business whilst delivering the commercial objectives.

However, as global data legislation has broadened, control environments have evolved in parallel and become increasingly complex. Controls to address new legislation are often layered upon (not integrated with) existing processes. This leads to control environments that are simultaneously siloed and duplicative, poorly embedded with inconsistent outcomes, and expensive to maintain and operate.

Some of the more common issues in the market include:

- A lack of understanding of what data is held, and where, within the organization, including unstructured data. As noted in the previous section, a well-maintained and comprehensive data catalogue that allows the business to understand its data asset is key to effective application of controls. Such an asset should have been implemented as a consequence of regulatory initiatives like BCBS 239 – not just data privacy regulations.
- Siloed management of consent and servicing of data subject access requests
- Poorly maintained, incomplete and inaccurate Records of Processing
- Cumbersome impact assessments that take months to complete
- Lack of granularity in privacy controls, blocking entire systems from certain processing use cases for a handful of attributes.

As this section sets out, a number of key considerations must be addressed to tackle these issues in a more holistic manner:

1. Streamlining the controls by removing duplication or otherwise overlaps in control function
2. Codifying the rules to improve consistency of outcomes and support automation
3. Introducing automation in the decision-making process
4. Leveraging the advantages of knowledge graphs to understand and manage the regulatory horizon
5. Doubling down on uplifting the ‘defensive’ elements of data literacy.

Improving control integration and interoperability

It is only possible to manage risks related to data privacy and protection under two conditions: having a trusted method to assess and make decisions on data sharing and access; and having a robust, discovered view of where the risk lies.

The Data Protection Impact Assessment (DPIA) and Records of Processing (ROP) are well known control processes that many FSOs implement to meet these key privacy control objectives. Implementing a DPIA provides a mechanism whereby intended data sharing / usage is effectively and efficiently vetted for risks, while the ROP facilitates knowledge of what sensitive data is processed where and for what purpose.

Beyond these two cornerstone controls are others that provide for, to name a select few:

- The management and servicing of data subject requests (DSRs)
- The gathering and management of consent
- The issuance of privacy notice
- The governance and monitoring of third parties with respect to data privacy and protection
- The management of privacy and protection breaches.

That the DPIA and ROP can effectively touch upon all of these puts these two controls at the very heart of all privacy implementations. Removing costs, ensuring successful embedding, and mitigation of data-related risk, involves a maturation journey that integrates and automates the operating model around these two controls. Effectiveness at scale requires not only improvements in their design and operation in isolation, but also in how their interrelations with the wider control suite are built in.

Essentially this involves a need to move away from controlling for risks via isolated control processes, to controlling via an integrated data / metadata architecture.

For instance, that legal basis and consent are necessary considerations in an impact assessment implicitly suggests the need to (a) ensure a globally consolidated record of legal basis and consent exists and (b) build an impact assessment that can readily query this. Furthermore, that existing data processes are associated with legal basis and/or consent implicitly suggests an opportunity for integrating these concepts with the ROP to further help streamline the control environment.

Additionally, looking more broadly, data privacy and protection implementations fundamentally require knowledge and management of what data is processed and where. The implication is that scalability is dependent on integration with data management constructs such as the data taxonomy and lineage, as covered in the previous section.

In consideration of the need to reduce the overall control burden therefore, the challenge of streamlining the privacy and protection control environment is primarily one of improving the interoperability and integration of the controls. There are opportunities to do so both between privacy and protection controls and with controls / and processes that typically exist outside of privacy and protection, in for example data management or even cyber security.

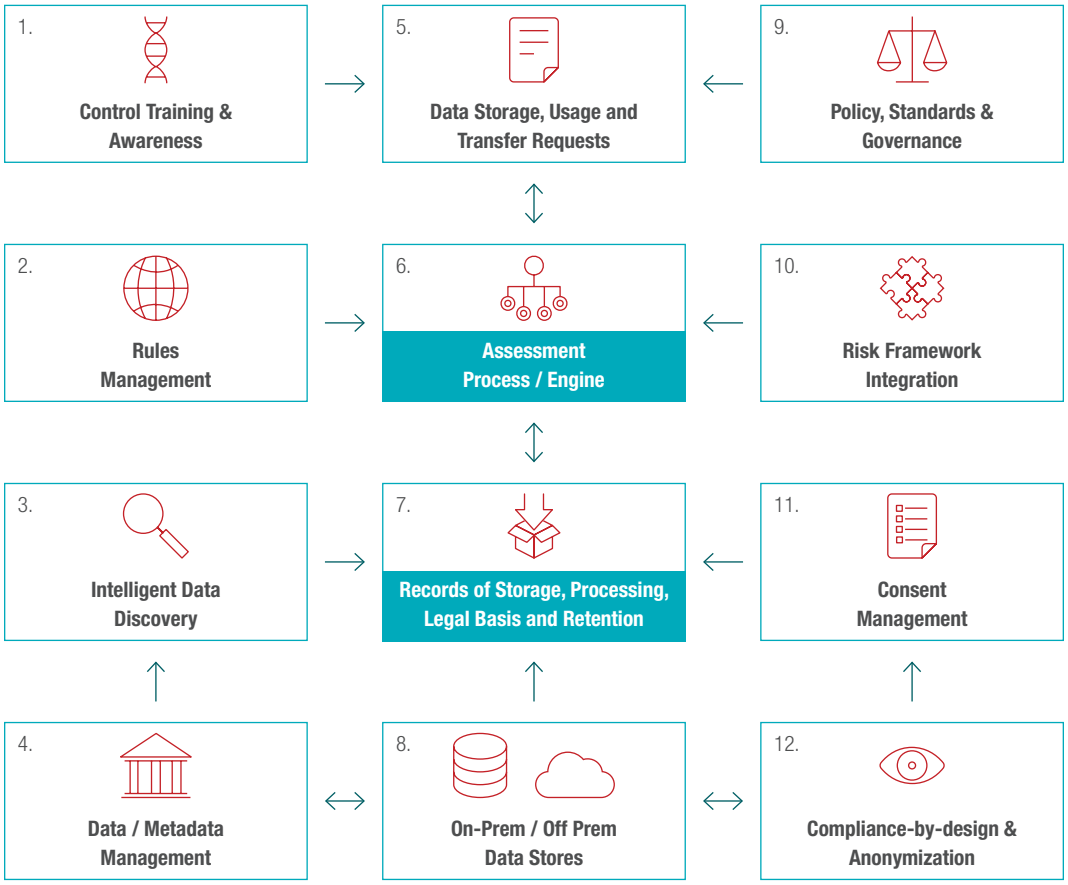

Figure 2: Integrated Control Model

A rules-based approach

While privacy and protection controls should be in theory symbiotic with an operational or strategic process, the practical reality is that they are still often regarded as a tax on the organization, impeding the efficiency of operations and the ability to innovate. This perception is in part because the controls are frequently set up as distinct processes rather than integrated into the core business processes – thereby creating unnecessary operational overheads.

Even if controls are streamlined and integrated, the ever-increasing variety of legislation only serves to make control processes more complicated due to the tight coupling of data controls to the laws that they serve. This results in high costs of adherence to new legislation which introduces changes from updates to control processes that must be redesigned, reimplemented, and retrained through to entirely new controls. There is therefore a pressing need to create a control environment that is accommodative of additional regulation at lower marginal cost.

The required insight is that, despite the multitude of guises that data protection laws manifest, they all seek to limit the collection, storage, transfer, access or processing of specific types of data. Controlling risk is therefore fundamentally a case of being able to control what is done (or required to be done), to what data, where, and for what purpose.

Understanding this allows for an approach that expresses all data legislation as structured rules governing the what, why and where of data on top of which a single control process can be built. This effectively decouples the process from its informational content, with the practical responsibility of rules maintenance shifting to the second line. Accommodating new legislation (or changes to existing legislation) is then a case of expressing a new set of rules with no need to amend the control process.

In this manner, a regulatory controls implementation can be easily adapted to service the requirements not only of other localized privacy laws but also other data protection legislation such as data residency and sovereignty.

Automating control processes

Impact assessments are typically highly manual processes requiring lengthy discourse between legal experts, control owners, and data requesters in gathering and processing the information necessary to make data protection decisions.

While this is necessary when legal text requires interpretation, codification of legal text into structured rules allows for the assessment itself to be conducted largely by automation and a rules engine. Initiatives led by regulatory bodies or third-party providers to digitize their rulebooks in machine readable form or even into code support this approach and simplify maintenance of the controls through updates to the codified legislation.

Automating the decision-making process in this manner has several benefits:

- Reduces the time taken to process a data access, transfer or usage request thereby improving the ability to scale with demand
- Improves the consistency of decision-making
- Improves the organization's ability to systematically learn and improve from historic cases
- Enables traceability and auditability of decisions
- Facilitates the impact analysis of changes in or introduction of new legislation on historic cases

It should be emphasized that data privacy and protection decisions are fundamentally subjective and risk-based in nature, and hence there always needs to be an element of human involvement – so full automation is an unrealistic, and indeed unwanted, goal.

However, automation does enable standard/common or otherwise low-risk requests to be processed automatically, with risk owners only involved in the sign-off on decisions. Only complex cases requiring bespoke judgement would require significant human involvement, and even this can be gradually reduced over time by introducing a machine learning element to the solution.

It should also be reiterated that the ability to operationalize any automation is reliant upon the deep understanding of the data across the organization's estate. Do we know what data there is and where it resides (discovery)? Have we tagged and classified the data (metadata)? If the organization has this understanding, together with the codified regulations, then automation and streamlined assessment workflows can be implemented.

Knowledge graph

Building solid data management foundations as articulated in the previous section allow the identification and understanding of the data being used – and to achieve these goals at scale. However, even in relatively modern data architecture designs, the movement of the data is considerable and adds complexity in relation to the limitations or conditions for data access must be applied.

Related to this, two further operational challenges present themselves:

- From the business perspective, as data driven initiatives gather pace, greater demand is asked of controls that assess the risk of data usage and sharing. The burden of manual execution becomes too great and creates a barrier to the swift access of data and innovation. Ideally, a solution to automate the controls will be explored to minimize control overheads as well as increase the speed with which the business can leverage its data assets.

- Regulatory and policy change (not just limited to data privacy), can be significant in the impacts that they have to the business in terms of the need to enable technical changes, updates to processes and the rollout of training. Significant manual time and effort is spent in the assessment, review and determination of targeted changes. The complexity and divergence of regulations, particularly for data privacy across multiple jurisdictions, is likely to increase as country-specific regulations are published. This, coupled with the desire for each organization to maximize the potential value of its data assets, means that a 'one-size fits all' approach that defaults to the most conservative regulation is neither feasible in the long term, nor optimal.

Knowledge graphs can be deployed to address these challenges. Utilizing captured metadata, a semantic, graph-based meta model – ideally linked to the wider enterprise data model – captures the information embedded in different data legislation/regulation.

Maintaining the rules in a graph-augmented risk and control framework allows for rationalization within and between different policy versions and different policy types. The graph could also be augmented to hold the process, role and system scope, and to identify relationships between these dimensions

From this, the knowledge graph can be used to:

- Identify the delta as regulatory or policy change occurs, noting impacted technical systems, users and processes. This allows for enhanced change precision, minimizes wasted effort and reduces overall risk arising from the incorrect inclusion (or exclusion) of impacted parties / change

- Identify contradictory regulations or policies that could cause conflicts around compliance with data privacy regulations/policies

- Link metadata repositories, such as the Record of Processing and/or data lineage, in a navigable integrated manner

- The graph-based metamodel allows for faster and more accurate interrogation to identify restrictions and limitations to sharing and processing of the data. This can be achieved at greater pace (via business rules automation), in support of providing accelerated access and impact assessment across multiple regulations.

The final point will become more pertinent given that the proliferation of jurisdiction-specific data privacy regulation shows no sign of decelerating, and organizations – either expanding into new territories or already established globally – will need to be able to easily assess and respond to this changing landscape.

Raising defensive aspects of data literacy

Data literacy is a critical enabler for organizations wishing to leverage investment in analytics. Employees in data-driven organizations need to operate in a particular way, and data literacy is about improving the confidence and fluency with which people think about, use, and manage data.

These efforts have tended to focus primarily on improving analytics capabilities and maximizing the potential value from the data asset. Defensive data literacy requires that programs to train and champion the adoption of data as a commercial competitive advantage are run in tandem with education streams on the governance and compliance processes that ensure the correct controls are applied to data use.

CONCLUSION

Data democratization poses significant challenges for organizations, especially those with a global footprint, when it comes to ensuring complete coverage and efficiency of operation. Historically, attempts to address compliance challenges by layering additional controls upon existing controls has resulted in processes that are simultaneously siloed and duplicative in nature – reducing operational efficiency.

We have covered the challenges associated with meeting regulatory requirements and principles in the context of leveraging the foundations of a strong data management capability and embedding these solutions within the control environment. The ability to meet these requirements and principles are increasingly pertinent given the need to adapt to the changing legislature and growth in data usage.

At its core, the challenge is one of controlling what data is involved, where, and for what purpose. Therefore, getting the fundamental 'Discover, Classify, Control' steps of data management right remains a crucial consideration. Firming up these data management foundations helps build scalability into the control environment.

Introducing automation in the control environment is key to reducing the cost and operational burden of controls. Approaches presented above have centered around rules-based automation – either in the creation of data assets that enable improved understanding of an organization's information, or in the execution of the controls themselves.

In conclusion, with market, technology and data use cases evolving rapidly, it is not adequate to focus on solving today's data storage and access control problems without also considering tomorrow's challenges. Data privacy and protection requirements are constantly evolving, and organizations need to adapt at pace to achieve and maintain a competitive advantage. An automated, rules-based control environment supported by solid data management foundations is a critical differentiator for businesses wanting to achieve scalable but ubiquitous compliance.

REFERENCES

1. londonlovesbusiness.com. 2022. GDPR fines hit nearly €100 million in H1 2022. 11 Jul. https://londonlovesbusiness.com/gdpr-fines-hit-nearly-e100-million-in-h1-2022/

2. legaljobs.io. 2022. 10 Eye-Opening GDPR Statistics. https://legaljobs.io/blog/gdpr-statistics/

AUTHORS

Stephen Brown, Managing Principal
stephen.brown@capco.com

Alvin Tan, Principal Consultant
alvin.tan@capco.com

CONTRIBUTORS

Gima Nwana, Senior Consultant
Charbel Khalil, Associate

ABOUT CAPCO

Capco, a Wipro company, is a global technology and management consultancy specializing in driving digital transformation in the financial services industry. With a growing client portfolio comprising of over 100 global organizations, Capco operates at the intersection of business and technology by combining innovative thinking with unrivalled industry knowledge to deliver end-to-end data-driven solutions and fast-track digital initiatives for banking and payments, capital markets, wealth and asset management, insurance, and the energy sector. Capco's cutting-edge ingenuity is brought to life through its Innovation Labs and award-winning Be Yourself At Work culture and diverse talent.

To learn more, visit www.capco.com or follow us on Twitter, Facebook, YouTube, LinkedIn Instagram, and Xing.

WORLDWIDE OFFICES

APAC
Bangalore
Bangkok
Gurgaon
Hong Kong
Kuala Lumpur
Mumbai
Pune
Singapore

EUROPE
Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
London
Munich
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Hartford
Houston
New York
Orlando
Toronto
Tysons Corner
Washington, DC

SOUTH AMERICA
São Paulo

WWW.CAPCO.COM

CAPCO
a wipro company

JN_4392