COM-PACE: Compliance-aware Cloud Application Engineering using Blockchain

xxx xxxx, *Senior Memberr, IEEE,* yyyy yyyy, *Senior Member, IEEE,* zzzz zzzz, *Fellow, IEEE,* yyyy yyyy, *Senior Member, IEEE,* and zzzz zzzz, *Fellow, IEEE*

Abstract—The necessity and obligation (like in COVID19 Pandemic) has escalated the dependence on the online services (from government, superstores, entertainment), often hosted over external Cloud computing infrastructure. The users of these services interact with a web interface rather than the larger distributed service provisioning chain that often involves an interlinked group of providers. The data and identity of users are often entrusted upon the service provider who may share it (or have automatic sharing agreement) with back-end services (such as advertising, analytics). To resolve this problem, General Data Protection Regulation (GDPR) has necessities the evolution of a compliance-conscious cloud application engineering that can provide an end-to-end solution for fair, transparent and lawful usage of users personal data. The existing state-of-the-art Cloud solutions and available SDKs have been concerning at infrastructure level rather then at compliance level. Taking a leap ahead, we propose a vision in the form of Compliance-aware Cloud Application Engineering (COM-PACE).This article provides an overview of key compliance operations and the perceived programming challenges for the realisation of these operations in current cloud infrastructure.

Index Terms—General Data Protection Regulation (GDPR), Blockchain, Cloud Computing, Security and Privacy, Smart Contracts.

1 Introduction

With the increasing demand of externally hosted services (from government, finance, entertainment), often hosted over Cloud computing infrastructure, there is a realisation that on-line electronic services can involve an interlinked set of providers. Users of these services only interact with a Web interface rather than the larger, distributed service ecosystem. They often endow (or entrust) their data and identity without perceiving that the service provider may share their data (or a portion of it) with several back-end services (Cloud hosted analytics, advertisers). While this has been a problem in the past, it will be considerably aggravated by the escalation of internet-connected devices in recent times. To overcome this, the General Data Protection Regulation (GDPR) is being implemented to ensure that non-expert users can make knowledgeable decisions about their privacy and thereby give 'informed consent' to use, store, share and reprocess their personal data. However, there are several challenges to facilitate this, both for individuals (data owners) who need to provide consent and for data controllers who need to obtain it.

One of the major challenges in the above context is the confusion of three terms, i.e., *Security*, *Privacy*, and *Compliance*. In general, these are inter-related but they have some distinct functionality that make them different in nature and technicality. *Security* refers to the freedom (resilience) from potential harm or damage (such as disruption or misdirection of services) caused by others (such as attackers, malware, virus, etc). *Privacy* relates to any entity or an information that is secluded from an individual or a group. It is linked directly to the sensitivity of data as any information that is private to an individual tends to be sensitive or personal. *Compliance* refers to as an *act of obeying*, i.e., any conduct that is based or bounded on (by) a specific rule, policy, order, or request. In other words, compliance signifies the conformance to a rule or guidelines (like a standard, legislation or law). Compliance is based on different principles, for example, the GDPR sets out seven key principles, Lawfulness, Fairness and transparency, Purpose limitation, Data minimisation, Accuracy, Storage limitation, Integrity and confidentiality (security), and Accountability. This reflects that Security and Privacy can be considered as individual entities but they are a few of the inherent principles that lead to compliance.

The top cloud providers (Amazon Web Services, Google Cloud, Microsoft Azure, IBM Bluemix, etc) either do not provide or provide limited support (restricted to security and limited privacy) for compliance adherence (specifically for GDPR). These cloud providers provide ready to use stacks (Serverless computing, Function-as-a-Service, CloudFormation, etc) but there key focus remains on the infrastructure level rather than compliance level. For instance, Function-as-a-Service provides a serverless platform (AWS Lambda, Google Cloud Functions, Microsoft Azure Functions, etc) where users can deploy, run and manage their applications without worrying about the infrastructure complexity and management (which is handled by cloud providers). Likewise, AWS CloudFormation provides the platform where you can code your infrastructure from scratch using the CloudFormation template, test it locally or at Amazon S3, use APIs/browser console/command line tools to create your own stack, and finally CloudFormation provisions and configures the stack as provided on your template. At infrastructure level, the cloud providers are

- *xxxx, yyyy, zzzz are with the School of Computing, Newcastle University, United Kingdom, NE4 5TG. E-mail: xxxx@ymail.com*
- *xxxx, yyyy, zzzz are with the School of Computer Science and Informatics, Cardiff University, United Kingdom, CF24 3AA. E-mail: xxxx@ymail.com*

ameliorating horizontally as well as vertically with respect to speed, scale and quality of service. But, the biggest threat, i.e., data breaches or loss of sensitive/personal data still cause distress among the organisations and users while using cloud infrastructure.

Let us consider an example of Amazon S3 bucket which backups the personal data related to users (let us say for an online pharmacy). Initially, all the permissions were set up properly and one can access the data remotely from any location. As the time passes, more and more data is added and at some instance you stop inspecting the data. Even, the original permissions were not verified from a long time. Now, what happens if this personnel healthcare data gets leaked? *Who accessed this personal data*? This question remains unanswered because the cloud providers basic definition is *'any authenticated user'*, i.e., anyone having valid account (like AWS IAM account), not even specific to the organisation. The AWS Identity and Access Management (IAM) allow you to create AWS user groups and manage their permissions related to AWS resources. So, any valid or authenticated user in the group created using AWS IAM capability can access the data. This goes strongly against the GDPR compliance principles (for example, there is a violation of purpose limitation, transparency, confidentiality, etc)

Although, the cloud providers have made significant mitigation's regarding the data leakage challenges but still there is a long pathway to go from here. The liabilities of any unauthorised access or usage of personal data can be huge. Specifically talking about GDPR, any violation of data privacy guidelines (scope is wide) can end up with huge amount of penalties. The GDPR necessities the organisations to report any kind of data breaches withing 72 hrs or hefty fines (4% of annual global turnover or 20 million Euros) are applicable [1]. For example, Facebook had to pay $5 billion due to *2018 Cambridge Analytica scandal* wherein 50 million profiles were accessed to target advertisement during the 2016 presidential election campaign [2]. In another instance, major penalty decisions are looming over Facebook under relevant GDPR regulations after the personal data of 50 million users was exposed in September 2018 due to a vulnerability in 'view as' feature [3].

One of the biggest questions for cloud providers is to understand the sensitivity of the data entrusted upon it by the users or the organisations. *What data is labeled as personal or sensitive while provisioning the cloud services in compliance-aware environment*? Whenever, there is a shared premises for data processing and storage, there lies the risk of data leakage that must be monitored continuously. *What are the conditions and applicable solutions for compliance monitoring*? Another problem lies with the compliance guidelines across different geographic locations as cloud data is stored at different locations. *How to understand the geographic data privacy regulations and apply them while monitoring the compliance data in a multi-tier environment*? It is quite hard to understand the applicable laws and monitor the data flows across different geographic perimeters. *How to verify the compliance and then ensure the 'right to be informed' clause through compliance enforcement*? The users or employees of an organisation often entrust there privacy commitments with the organisation itself but when the organisation relies on cloud providers providing infrastructure across different geo-locations then different privacy and compliance conditions arise. *How to tailor a shared agreement or responsibility related to compliance provisioning, monitoring and verification in a controller-processor scenario*? Summarising the above questions, the real concern is, *How do I create, monitor and then enforce compliance policies (through provisioning, monitoring and verification) that prevent breaches on the cloud*?

2 CE for hosting Services in the Multi-Cloud Environment

In this article, we present a snapshot of a tailored compliance aware solution in the form of *compliance-aware cloud application engineering* (in short compliance engineering (CE)). Fig. 1 shows the overview of proposed CE architecture in a distributed service ecosystem. To help understand this landscape, we characterize the CE in a multi-layered compliance-aware service stack (Fig. 1a). This architecture is build up over a traditional Infrastructure as a Service (IaaS) layer comprising of multiple cloud service providers which are managed through virtualisation (virtual machines or containers). Next comes the Platform as a Service (PaaS) layer comprising of serverless computing, database services, etc. Now, our work focus on how we make the scope of existing IaaS and PaaS wider to make it compliance aware through our new CE layer.

The application architecture decides how, when, and which compliance operations should be executed in the proposed architecture. The deployment of such compliance-aware architecture for providing services in the multi-cloud environment is very complex and challenging as cloud application composition involves intertwined inter-operational dependencies among the heterogeneous resources (software, hardware, VM/container). Fig. 1b depicts the high level architecture of the compliance application comprising different software resource layers including data layer, application logic, and compliance engine. Here, the different compliance operations are programmed to coordinate and control the application and compliance resources (at run time and design time) required to adhere to the compliance enforcement and usability. To follow our CE architecture, the cloud application engineers and deployment teams have to follow three inter-related steps and programming operations (shown in Fig. 1c) discussed below.

Provisioning compliance (at design and runtime): Here, the application owner analyses the user trust requirements or an organizations software resources to realise the required compliance requirements (including the data transfer constraint's) according to the applicable data protection regulations. After this, the compatible hardware resources are selected for instantiating the compliance-aware trust services and configuring them to handle the interoperability and communication with other software resources in the multi-tier web application. The amalgamation of compliance server or manager with the database server is a visible example in the Fig. 1b. For initiating the compliance provisioning, the data purpose and data usage contracts are established between the application owner, i.e., the data controller and the user or the customer, i.e., the data owner.

Monitoring compliance (runtime): The monitoring of operations performed on the data and performance metrics

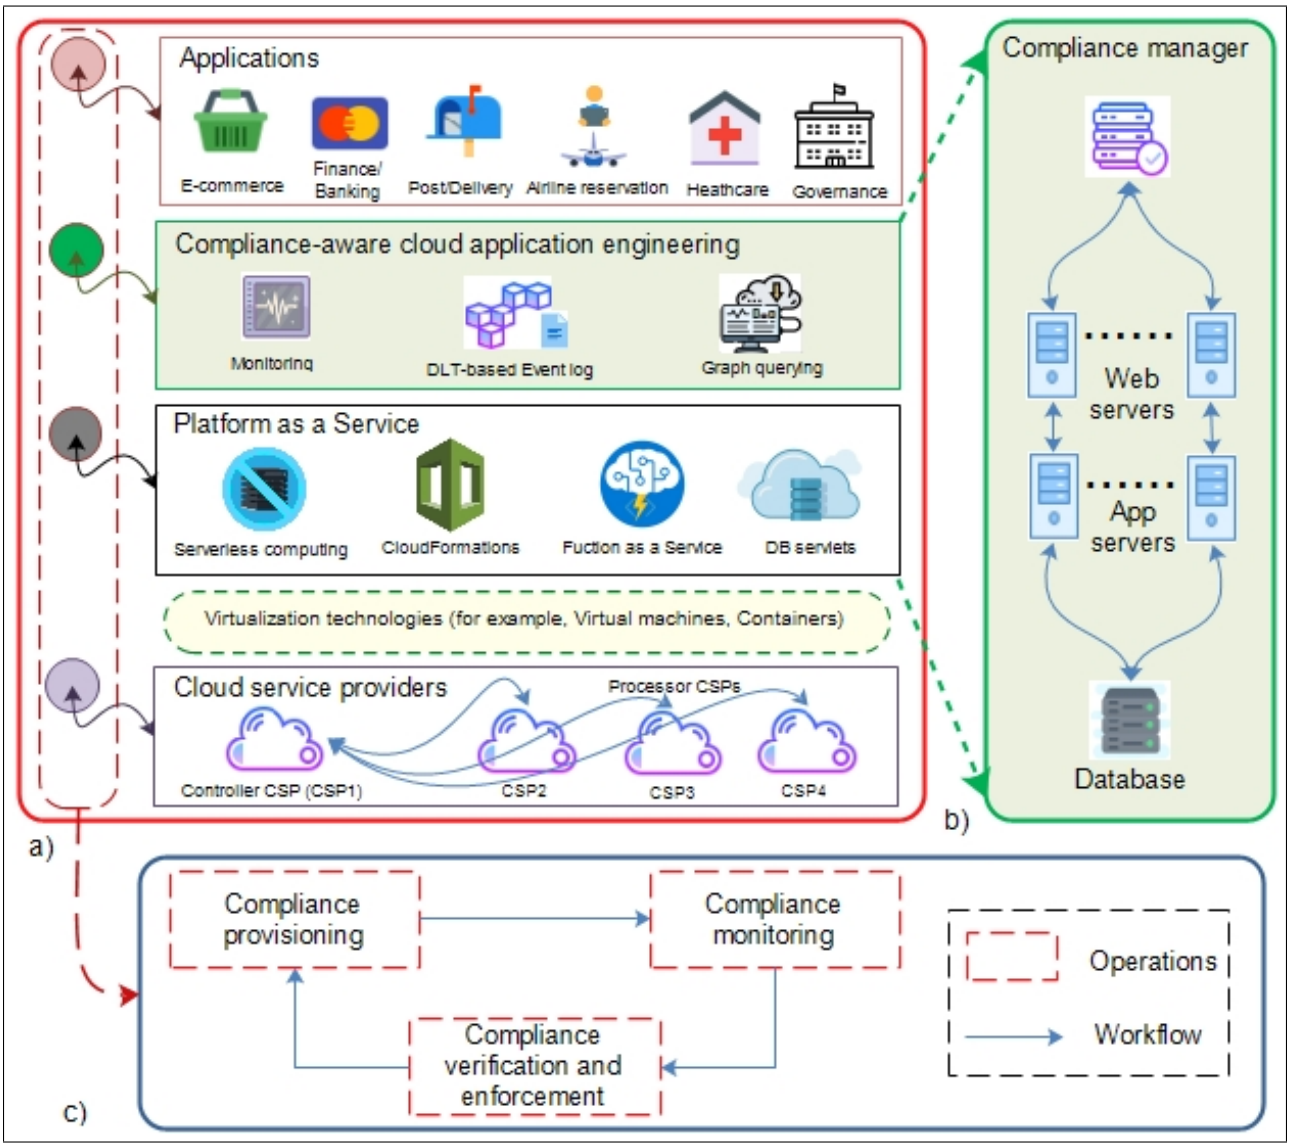

Fig. 1: Overview of compliance-aware cloud application engineering. a) Interlinked compliance-aware service stack that provides a layered architecture characterized on the basis of different attributes and granularity. b) High-level architecture of the layered cloud application engineering comprising of a compliance manager, web servers, application servers, and database server. c) Abstract view of the life cycle of the compliance engineering operations. (CSP: cloud service provider, DLT: distributed ledge technology)

related to the underlying cloud application helps to track the events which can be further correlated to GDPR violations (such as data leakage for analytics or advertising). This is attributed to the event information generated by the resources deployed for running the cloud application (like location) that can be used to understand the applicable data protection rules and relate it to the possible violation event recorded by the monitoring engine. A monitoring executor can initiate the data operations recording and subsequent submission to the blockchain for further analysis.

Verifying and enforcing compliance (runtime): Based on the compliance contract, the blockchain manager verifies the compliance through the query executor that analyses the behaviour of the events in line with the GDPR regulations. On verification, an event (such as data transfer, disclosure of personal data , profiling, processing of personal data for advertising, etc) can trigger the violation alert and thereafter data controller initiate corrective actions (such as reporting the violation to the regulation authorities) possibly without disturbing the runtime system.

2.1 Compliance Provisioning

Cloud SDKs provide a contemporary way of hosting web application components and provision data and services. However, the current APIs available for handling risk, governance, and compliance in these Cloud SDKs are not fully capable (or inconsistent) of provisioning compliance to the extent it is required. Some of the current cloud APIs (like Cloud Elements [4], Microsoft Azure [5], [6], AWS compliance programs [7], RedLock [8]) enable secure access to user assets (or personal data), but without adhering or limited adherence to the GDPR compliance principles. For

example, Cloud Elements [4] is an API integration Platform as a service (PaaS) hosted on Amazon AWS that provides stringent security provisions and practices, but, without any GDPR compliance control and enforcement. Although, Cloud Elements utilise and integrates different security solutions (for intrusion detection, file integrity, multi-factor authentication, etc), but it fails to provide a multi-level trusted environment for compliance provisioning. In such a solution, even though stringent security practices are adopted, but still it cannot track the data usage audit trail and operations that happened on the data. To overcome some of the above issues, the Cloud Elements Veracode application security program tries to provide a unified platform for protecting sensitive information (according to GDPR regulations) by using 256 bit AES encryption scheme and normalising APIs to connect to endpoints.

Microsoft Azure provide built-in compliance tools, but once the services have been provisioned at the provider end, the entire responsibility of operating the security and privacy policies rely on the user (who often is not fully trained to do so) [5], [6]. Amazon also provides a shared responsibility-based compliance cloud APIs wherein the role of management, operation and verification of security and privacy policies lie with the user and AWS handles only the overall operational chores [9]. For example, in Amazon EC2 service, the security configuration and management (including the configuration of the AWS-provided firewall configuration) is handled by the customer or user when they deploy Amazon EC2 instance. In the case of abstracted services (like Amazon S3 and Amazon DynamoDB), the customer is responsible for the data management (including encryption chores), asset classification, and application of relevant permissions [7]. Palo Alto Networks cloud-based RedLock service provides automatic redressal and compliance reporting along with ease of control in a multi-public cloud [8]. RedLock uses the APIs of the major public CSPs (like Amazon Web Services, Microsoft Azure and Google Cloud Safely) to provide an agent-less multi-public Cloud PaaS IaaS security environments for handling sensitive data. But, it is concerned only at network level and not at the application level.

For a fully compliance-aware environment, the need of end-to-end compliance provisioning is the first step to move. The current cloud APIs should be normalized to connect to many endpoints (like Overleaf, Dropbox, etc.). Compliance provisioning should enable the connection with the endpoint and thereafter the data can be streamed directly to the user or customers's application. During this process, any pass-through data from services should not be stored and the entire end-to-end transmission should occurs through HTTPS (supported with highest quality ciphers). Finally, the data stored at the endpoints is encrypted using the best possible encryption scheme (such as 256 bit AES encryption scheme). Compliance provisioning act as the backbone for delivering security and privacy solutions throughout the application life cycle. It should ensure the continuous scanning of all the application and components (build or purchased) even covering all the frameworks, application types and so on.

2.2 Compliance Monitoring

Compliance monitoring is mainly responsible for providing affirmation that an amenable framework is being adhered to as a watchdog for unwarranted operations or event, and acts as an autonomous process, operating in the second line of defence [10]. The base of compliance monitoring can be assumed from the process execution events that follow up. Most common frameworks for monitoring track complex events and submit them to the monitoring tier and significant events can be accumulated from real low-level execution events [11]. The monitoring tier ensures obedience to the adhered rules. Concerning the stakeholder's requirements, visualization of results is carried out in the reporting tier, which in turn, attains input from the monitoring tier.

Currently, various monitoring frameworks (like docker stat [12], cAdvisor [13], DataDog [14], Amazon cloud watch [15], CLAMS [16]) are available to observe the applications running in the cloud. But, these monitoring frameworks are either cloud provider-oriented (Microsoft Azure Fabric Controller) or virtualization architecture-oriented (cAdvisor) [17] and hence fail to meet the monitoring requirements in a complex multiple cloud environment. Several studies [16], [17] have been carried out, all of which concentrated on the comprehensive performance-based monitoring in the Cloud. But, to be specific for GDPR compliance, the exact data processing event during individual stages of the process cannot be portrayed by the overall metrics. There remains a lack in the GDPR metrics and an acute need for an intelligent monitoring framework. An extensive investigation of both the real-time monitoring overhead and the framework scalability is also required. Hence, monitoring what and monitoring how appears to be the primary challenges of multi-cloud event monitoring framework where the event logs have to be stored in a blockchain.

To combat these challenges, we put forth the proposal for an extensive real-time compliance monitoring framework that can be used to monitor the processing of personal data (in line with GDPR) in the multi-cloud systems. Using the daemon process and log analysis, the performance data is obtained by the framework in real-time. The dimensions of compliance monitoring are elaborated below.

2.2.1 Monitoring Granularity

A wide variety of technologies (virtual machines, Docker Containers) have been used to increase the stack of elements that must be managed for application creation, including the use of containers to run the software, Web servers or big-data processing. Although typically only hardware and software structuring components (servers, databases or proxies) needs to be controlled, monitoring at the lower rates, i.e. cloud systems, microservices and APIs, even internally used methods or functions, is progressively needed. The purpose of this CE architecture is to provide an automated management framework for the different layers used in applications decomposed in microservice architecture and container clusters.

Not only on a single host but in many container clusters, a container-based application can be deployed. There are several nodes for each container cluster (hosts) and there are several containers for each node. Output management

data can be obtained from various code layers (e.g. node layer, database layer, and server layer) for applications implemented in container-based environments. The system seeks to mitigate that issue by introducing a multi-layer monitoring system for applications that are broken up in and deployed in containers in the multi-cloud environment (hardware metrics from operational monitoring, software metrics by monitoring server processes and database processes, besides internal metrics such as method latency or processing rates of an API call).

2.2.2 Monitoring Topology

The event capturing defined by the metrics is performed through monitoring agents, condensed in the conceptual representation of SmartAgent (SA). These are deployed at each CSP container and are responsible for monitoring all the events happening in the concerned container. In CE architecture, we have considered six types of agents connected to SA. The agents for monitoring the metrics related to containers, applications, file, user access, an external device, and network are considered in CE architecture. Here, the biggest concern is related to the selection of appropriated monitoring agent in the considered topology to make sure the compliance-related event is captured successfully. Another concern involves the selection of suitable compliance-related events to be monitored across the entire topology with a trade-off between monitoring overhead, response time, scalability and compliance coverage.

The challenges for monitoring frameworks in multi-cloud environment are discussed below.

Volume of Events and Alert Overload: The effect of the increasing resources of microservices, cloud, and virtualization on the monitoring needs is not always determined by a lot of organizations from beforehand, and thus, segregating the unexpected and expected can be difficult for containers which are ephemeral and dynamic. Conventional monitoring systems may have difficulty coping or may run out of event log resource allocation due to an overwhelming increase in event activity due to the easy scaling of resources by application owners. To tackle the problem with the tooling expansion, there lies an incoherence, where the concept of legacy environment is not familiar to modern tools, while containerized environments are not familiar to the conventional systems of monitoring.

Monitoring Rule Complexity: Monitoring of operations can become arduous as complex changes in the environment are expected to be notified to the users by the tools. Disregarding many probable real events and consideration of false alerts by the monitoring agents can cause problems. This kind of situation occurs mostly when a lot of alerts arise due to the imprecise rules and large volume, which make it difficult to be monitored by the monitoring agents.

Architecture of Complex Systems: Problems arise for monitoring systems due to systems and networks with varying degrees of trust, as communication is necessary within and beyond the boundaries of the trust. Such problems may be solved by using a unique trusted configuration for the monitoring system. Besides, several monitoring systems can also be used which can communicate with each other for comparing the information attained. The installation of such a monitoring system depends on the level of sensitive data that is dealing with, and also on the engineers who install the monitoring system in a controlled environment.

Auditing Issue in Clouds: The strategies for auditing and monitoring for attaining the duties of compliance are not integrated into the present public offerings of IaaS. In case of deployment of the application in the public cloud, it is difficult to follow the specifications of storage and location-based processing since the cloud properties maintain the underlying details abstractly. The deployment of business applications is no longer possible due to the inability to fulfil the requirements of compliance as the cloud cannot be monitored [18]. This inability to satisfy the terms of compliance may cause large fines or annulment of the business permissions. Hence, IP administration and VM scale operations should be issued with the auditing logs and monitoring facilities by the IaaS clouds. To assure the fulfilment of the audit logs, stern logging conditions are required. Using resources of the cloud in a traceless way should not be allowed even when the administrative account is used for logging in. The service provider constraints should also be aimed to be fulfilled by the IaaS. Any IaaS should bear the abundance of regulatory restrictions and legislations, and this should be ensured by the service providers.

Application Migration in Multi Clouds: Another problem which arises is related to the active organization of the containers or applications of the CSPs. Fixing the procedure and timing for the migration of applications from different CSPs and determining the properties that impact the migrations are a few examples to quote. Application parts are managed and deployed in different ways, as different CSP's deal with the process. It is difficult to manage the application components overall as a whole entity as the CSP's maintain a heterogeneous nature. A major part is played by the monitoring in recognizing the timing of migration of particular applications or containers.

Data Aggregation: Finally, data accumulation or aggregation poses a difficulty, as this is used to check the number of events which have been collected. Also, the events in a container lead to creating log files, which is also overlooked by the existing data aggregation technologies.

2.3 Compliance Verification and Enforcement

Here, the data (event) logs can be queried to verify and enforce compliance using smart contracts. The smart contracts are used to digitally verify, or enforce the compliance as necessitated by the contract. This helps to verify credible transactions without third parties invention. There are two major challenges, 1) to select an appropriate blockchain platform, and 2) to select the events that should be queried or verified to avoid additional overhead.

2.3.1 Smart Contracts-based Enforcement

Figure 2 presents an abstract model for connecting the parties, involving user, cloud service providers and arbiter to blockchain to use the smart contracts supporting GDPR requirements. The model enables the audit trail of service

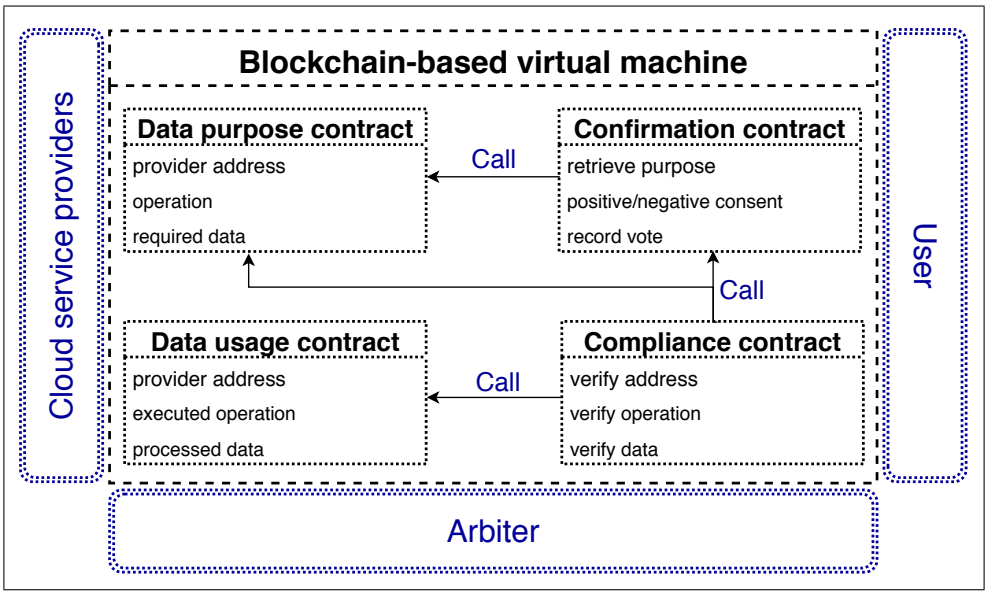

Fig. 2: The proposed smart contracts

providers that can have the roles of data controllers or data processors. It makes use of blockchain to record the operations (e.g. read, copy etc) carried out by providers on user data. Furthermore, the model checks whether the executed operations comply with GDPR or not. The blockchain-based virtual machine is an open blockchain platform (e.g. Ethereum virtual machine [19]) providing an environment for the parties to run smart contracts and create a blockchain network. The virtual machine involves four smart contracts that provide the basis for the verification of providers following a set of GDPR obligations. The smart contracts are data purpose contract, confirmation contract, data usage contract, and compliance contract.

Data purpose contract captures the purpose of data processing of cloud providers. The purpose of data processing can be specified with several typical operations (i.e. read, copy, transfer, profiling etc) carried out by controllers/ processors on persona data. The contract enables the providers to store their addresses (e.g. Ethereum accounts) and the operations that will be executed by them on user data in a blockchain. The activators of the contract's transactions are cloud providers. This contract realizes the Art. 30(1)(b) of GDPR under which the purpose of data processing and the address of service provider should be clarified in advance.

Confirmation contract enables users to give a positive or negative consent for data processing recorded in the blockchain by data purpose contract. The contract can contain two functions: one for retrieving the records containing the data processing purposes of providers; the other for sending the vote (accept/ reject) of users into a blockchain. The former permits users to retrieve the blockchain and access the purposes of data processing before sharing their personal data with cloud providers. Through the latter function, the user can specify whether the execution of an operation is allowed or not. The activators of such functions are users. The contract meets the Art. 6(1)(a) of GDPR under which data subject (user) has the right of giving consent for processing their personal data.

Data usage contract records all operations of providers carried out on personal data in a blockchain. Such operations are captured by the container on the provider's side to track the processes of the service provider on personal data. The contract can involve a function activated by container to store: provider address and executed operation. This information is sent to the blockchain to lay a basis for the verification of providers. The contract enforces providers to receive their users' consent before any data usage. Moreover, it enables users to track and be aware of the history of data movement among cloud providers. Such a capability supplying by the contract meets the Art. 15(2) and 20(2) of GDPR under which data subjects have the right to information about where their data are processing.

Compliance contract verifies the blockchains created by the aforementioned smart contracts to detect any GDPR violation. The contract is deployed and executed by the arbiter who is a trusted third party connecting to the blockchain virtual machine to report the cloud providers committing a GDPR breach. The following verification is undertaken through the contact to automatically clarify possible violators:

1) whether the addresses of service providers recorded by data usage contract conform to those recorded via data purpose contract or not;
2) whether the operations of each service provider recorded by data usage contract conform to those recorded via data purpose contract or not;
3) whether the operations of each provider recorded by data usage contract were already confirmed by the data subject or not.

ACKNOWLEDGEMENT

This work is supported by the Engineering and Physical Sciences Research Council (EPSRC) funded project PACE: Privacy-Aware Cloud Ecosystems (EP/R033293/1, EP/R033439/1).

REFERENCES

[1] G. Strawbridge. (2018) 5 examples of security breaches in 2018. [Online]. Available: https://www.metacompliance.com/blog/5-examples-of-security-breaches-in-2018/

[2] M. Snider and E. C. Baig. (2019) Facebook fined $5 billion by ftc, must update and adopt new privacy, security measures. [Online]. Available: https://eu.usatoday.com/story/tech/news/2019/07/24/facebook-pay-record-5-billion-fine-u-s-privacy-violations/1812499001/

[3] N. Lomas. (2020) First major gdpr decisions looming on twitter and facebook. [Online]. Available: https://techcrunch.com/2020/05/22/first-major-gdpr-decisions-looming-on-twitter-and-facebook/

[4] (2020) Cloud elements platform security and compliance. [Online]. Available: https://cloud-elements.com/security-compliance/

[5] R. Waggoner. (2017) Achieving trust and compliance in the microsoft azure cloud. [Online]. Available: http://blog.mycloudit.com/achieving-trust-and-compliance-in-the-microsoft-azure-cloud

[6] R. O'Leary and H. Singh. (2019) Microsoft azure is helping organisations manage regulatory challenges more effectively. [Online]. Available: https://azure.microsoft.com/en-gb/resources/azure-is-helping-organizations-manage-regulatory-challenges/

[7] (2017) Amazon web services: Risk and compliance. [Online]. Available: https://d0.awsstatic.com/whitepapers/compliance/AWS_Risk_and_Compliance_Whitepaper.pdf

[8] Continuous security and compliance for multi-cloud deployments by palo alto. [Online]. Available: https://www.digitalmarketplace.service.gov.uk/g-cloud/services/894784684094117

[9] Shared responsibility model. [Online]. Available: https://aws.amazon.com/compliance/shared-responsibility-model/

[10] D. Loreti, F. Chesani, A. Ciampolini, and P. Mello, "A distributed approach to compliance monitoring of business process event streams," *Future Generation Computer Systems*, vol. 82, pp. 104–118, 2018.

[11] L. T. Ly, S. Rinderle-Ma, D. Knuplesch, and P. Dadam, "Monitoring business process compliance using compliance rule graphs," in *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*. Springer, 2011, pp. 82–99.

[12] Docker. [Online]. Available: https://www.docker.com

[13] cadvisor (container advisor). [Online]. Available: https://github.com/google/cadvisor

[14] cadvisor (container advisor). [Online]. Available: https://github.com/google/cadvisor

[15] Amazon cloudwatch. [Online]. Available: https://aws.amazon.com/ cloudwatch

[16] K. Alhamazani, R. Ranjan, K. Mitra, P. P. Jayaraman, Z. Huang, L. Wang, and F. Rabhi, "Clams: Cross-layer multi-cloud application monitoring-as-a-service framework," in *2014 IEEE International Conference on Services Computing*. IEEE, 2014, pp. 283–290.

[17] A. Noor, D. N. Jha, K. Mitra, P. P. Jayaraman, A. Souza, R. Ranjan, and S. Dustdar, "A framework for monitoring microservice-oriented cloud applications in heterogeneous virtualization environments," in *2019 IEEE 12th International Conference on Cloud Computing (CLOUD)*. IEEE, 2019, pp. 156–163.

[18] P. Massonet, S. Naqvi, C. Ponsard, J. Latanicki, B. Rochwerger, and M. Villari, "A monitoring and audit logging architecture for data location compliance in federated cloud infrastructures," in *2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*. IEEE, 2011, pp. 1510–1517.

[19] (2020) Ethereum. [Online]. Available: https://www.ethereum.org/