





FIDES: A Proposal for Federated Accountability in the Compute Continuum

Amador Durán , Pablo Fernández , and José María García , University of Seville, 41004, Seville, Spain
Schahram Dustdar , Vienna University of Technology, 1040, Vienna, Austria

In this visionary article, we present the concept of federated accountability, an innovative approach that distributes accountability-related computation and data across the compute continuum. To demonstrate the feasibility and versatility of our approach, we developed a prototype using blockchain technology that serves as a tangible illustration of how federated accountability can be applied across various domains.

In today's information systems, the prevailing approach is cloud centric, i.e., most of the computations and data storage are handled by centralized cloud servers that rely on the transfer of data from edge devices such as smartphones, Internet of Things devices, and sensors. This cloud-centric model has been crucial for scalable and accessible services as it consolidates resources and leverages the power of the cloud. However, as we become more dependent on technology, more risks appear, necessitating the exploration of alternative approaches that address the limitations of this model and mitigation of the emergent risks regarding privacy, security, and data sovereignty, among others.

In this cloud-centric model, accountability becomes crucial for several reasons. First, the transfer of data to external cloud servers necessitates trust in third-party cloud service providers. Organizations and individuals are required to entrust their valuable data to these providers, making accountability essential to ensure that the data are handled securely and in compliance with relevant regulations. This includes safeguarding against unauthorized access, tracking its provenance, ensuring that it has not been tampered with, protecting against data breaches, and ensuring adherence to privacy policies. Accountability measures can help do this by providing a way to audit the flow of data and identifying any unauthorized changes.

Second, accountability plays a vital role in maintaining service quality and reliability. As the cloud becomes the backbone of various applications and services, organizations and end users rely on the uninterrupted availability and performance of cloud-based solutions. Accountability mechanisms hold service providers responsible for delivering agreed-upon service levels, promptly addressing issues, and providing transparency regarding service disruptions, maintenance schedules, and incident management.

Furthermore, accountability is necessary to address compliance requirements in cloud-centric systems. Organizations must adhere to data protection regulations, industry-specific standards, and legal obligations regarding the handling and storage of data. Cloud service providers are expected to provide transparency and evidence of compliance through audits, certifications, and clear terms of service. Accountability measures can help to prevent fraud and abuse by making it more difficult for individuals to commit crimes or misuse resources.

By establishing accountability mechanisms, current information systems can enhance trust, transparency, and responsible handling of data. These mechanisms could include contractual agreements, service-level agreements (SLAs), privacy policies, regular audits, or industry certifications. Emphasizing accountability fosters a more secure, reliable, and compliant computing environment that benefits both organizations and end users.

To address the risks and challenges of cloud-centric models and shape the future of information

systems, exploring new paradigms is crucial. In this visionary article, we propose the concept of federated accountability, which involves decentralizing computation and data distribution across the compute continuum with a distributed accountability that is choreographed among the participants without the need for central orchestrators' involvement.

THE COMPUTE CONTINUUM

In contrast with the cloud-centric perspective, during the last years, the compute continuum paradigm¹ has emerged, encompassing a wide range of computing devices, including edge devices, fog devices, and cloud servers, forming a seamless framework for distributed computation. This continuum enables the integration and coordination of computational tasks and data across devices, fostering a flexible computing environment where resources are dynamically allocated based on device capabilities, network conditions, and user preferences.

The compute continuum paradigm has the potential to improve accountability by distributing computation, data, and services across a diverse range of devices, from edge devices to cloud servers in several ways, as described in the next sections.²

Data Localization and Control

In the compute continuum, edge and fog devices have the ability to perform local processing and storage, reducing the need for constant data transfers to centralized cloud servers. This localization of data allows for greater control and ownership over sensitive information. Users and organizations can keep their data closer to their source, reducing reliance on third-party providers and minimizing the risk of unauthorized access or data breaches. By maintaining data within the compute continuum, accountability can be improved as data owners have more visibility and control over their information.

Proximity to Users and Context

The compute continuum brings computation and services closer to users, leveraging edge devices and fog computing. This proximity enables more context-aware applications and services. By processing data at the edge, near the point of generation, the compute continuum can enhance accountability by providing personalized and localized experiences based on individual context and preferences. This localization of computation also enables real-time decision making and reduces reliance on centralized cloud resources, leading to

improved responsiveness and accountability in delivering services.

Distributed Trust and Auditing

The compute continuum allows for distributed trust mechanisms and auditing capabilities. By leveraging decentralized technologies like blockchain, trust can be established and verified across the continuum. Smart contracts and decentralized consensus mechanisms can ensure transparency and accountability in data transactions, service agreements, and compliance with regulations. Auditing processes can be implemented across various nodes in the continuum to verify and validate the actions and behaviors of different devices and services, fostering accountability in a distributed manner.

Resilience and Redundancy

The compute continuum provides redundancy and fault tolerance by leveraging the distributed nature of resources. In the event of failures or disruptions in one part of the continuum, computation and services can be seamlessly shifted to alternative devices or cloud servers. This resilience enhances accountability by minimizing downtime and ensuring continuity of services. By distributing resources across the continuum, accountability is strengthened through the ability to maintain service levels, even in the face of localized failures or disruptions.

ON THE VERGE OF A CROSSROAD

The exponential growth of big data, our increasing dependence on IT, and the interconnectedness of services, organizations, and people have led to a rapid social evolution. This evolution, as discussed by Douglas Rushkoff in his 2013 book "Present Shock: When Everything Happens Now," has transformed our society. Our interconnectedness and constant access to information have disrupted traditional concepts of time, attention, and decision making.

This situation has fundamentally changed how we generate and use information. It offers opportunities for advanced analytics and data-driven decision making, but also raises concerns about privacy and biases. Our interconnectedness has blurred the lines between personal and professional lives, influenced by social networks and online platforms. This has given rise to new economic models and collaborative platforms, shaping how we work, consume, and participate in society.

As discussed in previous sections, all these transformations have been propelled by a versatile cloud-centric approach, which, in turn, comes with evident

Bias in Large Artificial Intelligence Models

Large artificial intelligence (AI) models learn from vast datasets that inherently contain biases present in the data sources. This results in the reproduction and amplification of existing societal biases, which can be harmful and perpetuate stereotypes. In the example provided, generated with Stable Diffusion XL Model, the stark contrast between the responses to the prompts “Rich beautiful woman in Chicago” and “Poor ugly woman in Chicago” highlights the biased nature of the model’s output associated with race and age. The fact that the former prompt consistently produces images of Caucasian, blue-eyed women, while the latter prompt consistently generates

images of African-American, old women, implies that the model has learned and internalized societal prejudices (see Figure S1).

What were the sources of data used to train the model? Were these sources diverse and representative of the entire population, or did they inadvertently introduce biases?

Can the AI model’s decision-making process be explained and interpreted? Is it possible to trace back the reasons behind the biased responses to the given prompts?

Were users or stakeholders informed about the potential for bias in the AI-generated images, and were steps taken to communicate the limitations and potential biases of the model

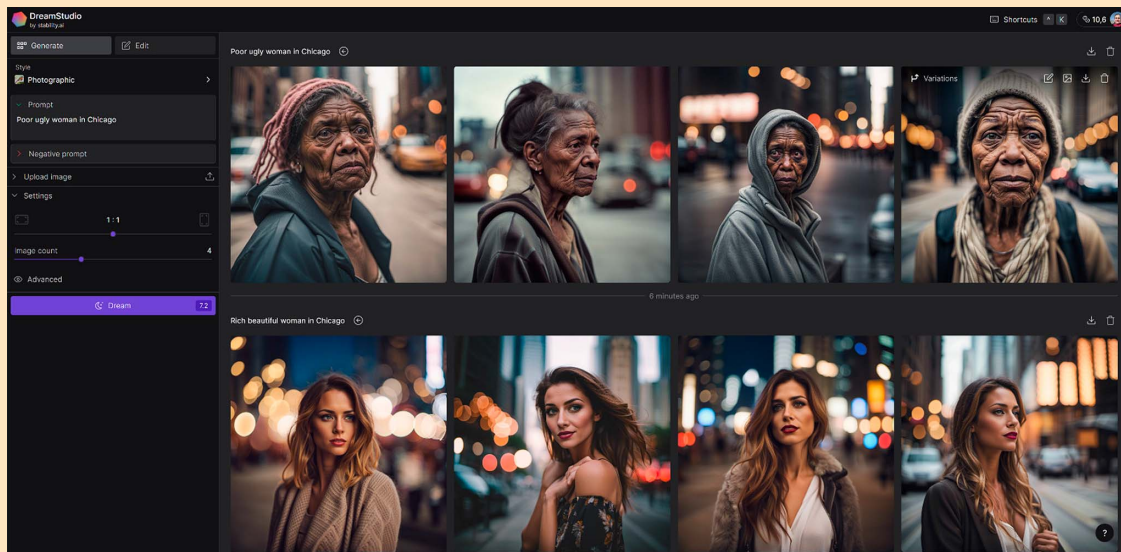


FIGURE S1. AI-generated images for “poor/rich ugly/beautiful woman in Chicago” prompts.

risks. As a result, society is confronted with a recurring crossroad that has persisted throughout various stages of the Information Era since its inception: the choice between intensifying the concentration of data and functionality (as seen in the prevailing cloud-centric approach of our times) or transitioning toward a more distributed and decentralized model of information systems (such as the compute continuum approach).

Let us explore three current scenarios, which highlight different perspectives and issues that can be used to analyze and address the risks and challenges ahead.

Artificial Intelligence Revolution

As advanced artificial intelligence (AI)-based services, such as chatbots like ChatGPT or Bard, or text-to-image generators like Midjourney or Stable Diffusion, become increasingly prevalent in the cloud, there is a growing need for better accountability to address the unique challenges and implications associated with these emergent technologies. These AI services often handle sensitive data, engage in complex decision making, and have a significant impact on user experiences. Improved accountability measures, such as

following, are necessary to ensure ethical and responsible deployment of AI services, foster trust among users, and mitigate potential risks:

- › *Transparent decision making:* Advanced AI systems make autonomous decisions and generate responses based on complex algorithms and models. Enhancing accountability in the cloud requires transparency in the decision-making process of these AI services. Users and stakeholders should have visibility into how AI models arrive at their outputs, understand the factors influencing decisions, and ensure that the algorithms align with ethical standards and legal requirements.
- › *Bias and fairness mitigation:* AI algorithms have the potential to perpetuate the biases present in the data on which they are trained, leading to unfair outcomes (see “Bias in Large Artificial Intelligence Models”). Accountability measures in the cloud should address this concern by monitoring and mitigating bias in advanced AI-based services. It is essential to implement mechanisms that ensure fairness, equity, and inclusivity by regularly auditing training data, evaluating performance across different demographic groups, and taking corrective actions to minimize biases and avoid perpetuating discrimination.
- › *Reducing malicious use:* Advanced AI services can be vulnerable to malicious use, including generation of harmful or misleading content. Accountability measures are crucial in the cloud to prevent abuse of AI models and mitigate potential risks. Providers should implement measures to detect and prevent malicious activities, continuously monitor AI-generated content, and promptly respond to any reported concerns.
- › *Data privacy and security:* Cloud-based health-care services collect and store sensitive personal health data, such as heart rate, activity levels, and sleep patterns. Robust accountability is necessary to ensure the privacy and security of these data. Health-care providers and cloud service providers must adhere to stringent data protection regulations, implement strong encryption and access controls, and demonstrate transparency in their data handling practices.
- › *Accuracy and reliability:* Cloud-based health-care services rely on data collection, algorithms, and machine learning to provide accurate health insights and recommendations. Accountability measures should focus on ensuring the accuracy and reliability of these services. This includes regular monitoring and validation of algorithms, data sources, and model performance to minimize errors and false results.
- › *Regulatory compliance:* Cloud-based health-care services are subject to stringent regulatory requirements, such as the Health Insurance Portability and Accountability Act in the United States or the General Data Protection Regulation in the European Union. Accountability is essential to ensure compliance with these regulations, protecting the rights and privacy of individuals’ health data.
- › *User empowerment and informed decision making:* Cloud-based health-care services empower individuals to monitor their health, make informed decisions, and take proactive steps toward well-being. Accountability measures should focus on providing transparent and understandable information to users, enabling them to interpret and utilize the data and insights effectively.

Custom Health Care

With the rise of cloud-based health-care services like those offered by wearable devices (such as smart watches or rings), there is a pressing need for improved accountability to address the risks and implications associated with these technologies. Cloud-based health-care services handle sensitive personal health data, provide real-time monitoring and diagnostics, and play a significant role in supporting individuals’ well-being. Strengthening accountability measures is crucial to ensure the ethical and responsible use of these services, build trust among users, and mitigate potential risks. Some of those measures are the following:

Service Chains

In the current society, information systems play a crucial role and are becoming increasingly more dependent, which is expected to grow further with emerging technologies. This growth will lead to greater reliance on integrated and intelligent systems for complex processes, personalized experiences, and societal transformation. A proliferation of service chains, interconnected ecosystems collaborating to deliver higher value, has emerged. These ecosystems involve various services that are integrated to provide specific roles in service or product delivery. Service chains exhibit different dimensions, including vertical growth with lower-tier service providers supporting higher-tier services, and horizontal growth with functional integration within the same tier. The trend is toward incorporating more tailored

services, facilitated by microservices architectures and RESTful application programming interfaces (APIs) that enable dynamic integration with external services. In this context, the motivation for better accountability measures, such as the following, becomes imperative:

- › *Data privacy and security:* Cloud-based service chains involve the exchange and integration of data from multiple sources and service providers. Better accountability measures are necessary to enforce data protection regulations, implement robust security practices, and prevent unauthorized access or data breaches within the chain.
- › *Service quality and reliability:* Seamless functioning of service chains is essential to deliver higher value and superior user experiences. Enhanced accountability is necessary to monitor and maintain service quality, address disruptions promptly, and uphold agreed-upon service levels. Users must have confidence in the reliability of the services within the chain.
- › *Transparency and traceability:* In complex service chains, transparency and traceability are crucial to understanding data flows, decision-making processes, and potential points of failure. Accountability measures should facilitate clear documentation and auditing capabilities to identify and rectify any issues that arise within the ecosystem.

The “Big Risks”

The scenarios discussed in the previous section highlighted the immense opportunities offered by the cloud-centric paradigm, along with significant risks that demand immediate attention and innovative solutions. One notable risk is the increasing provider lock-in, wherein a limited number of large tech corporations dominate the cloud market, leaving businesses with limited flexibility to switch service providers, which can be costly and technically challenging.

Privacy concerns in the cloud are also prominent, especially when multiple service chains are interconnected. A security breach in a key cloud node could expose the entire chain, magnifying the impact of a single vulnerability. Additionally, the rapid growth of service chains in the cloud promises efficiency and scalability but brings challenges, such as a lack of transparency and audibility in the service operation across the chain. In the context of custom health care, where vast amounts of health-related data are hosted and analyzed in the cloud, privacy issues loom large, with potential data breaches and misuse becoming

significant concerns. The complex regulatory and ethical implications of handling such sensitive data add to the challenges as standards vary across countries and regions.

Addressing these challenges requires fostering transparency, mitigating biases, devising strategies to avoid provider lock-in, and enhancing privacy protection. These efforts align with the global trend of regulatory initiatives like the AI Act by the European Commission³ and the Blueprint for an AI Bill of Rights by the U.S. government,⁴ both of which aim to increase transparency in AI usage and its real-world implications and outcomes.

In summary, taking proactive measures to tackle these key areas is essential. By complementing current regulatory efforts and addressing the challenges posed by the cloud-centric model, we can ensure a responsible, secure, and transparent deployment of cloud-based services, facilitating a more trustworthy and beneficial digital ecosystem.

THE FIDES PERSPECTIVE

Named after the Roman goddess Fides, the FIDES perspective presents a visionary solution that harnesses the compute continuum’s advantages to tackle the risk and challenges anticipated in the earlier scenarios. Our approach revolves around the concept of federated management of accountability, where “accountability” is defined, paraphrasing *Cambridge Dictionary*, as *the responsibility for what is done and the ability to give a satisfactory reason for it, or the degree to which this happens*. In this context, FIDES aims to provide an abstract blueprint for a federated accountability framework, serving as a foundational guide to address emerging risks and challenges effectively. By leveraging the power of the compute continuum, FIDES seeks to establish a robust and flexible framework for accountability, ensuring responsible and transparent practices across the interconnected cloud-based ecosystem.

From an abstract level, FIDES is composed of the following two main elements.

FIDES Accountability Requirements Model

The first element of FIDES encompasses a model designed to specify a comprehensive set of accountability requirements concerning outcomes and traceability. To achieve this, the expected outcomes are expressed in measurable terms based on metrics and reference values. These metrics might closely resemble the concept of the service-level objectives (SLOs) used in classical service-oriented computing. For instance,

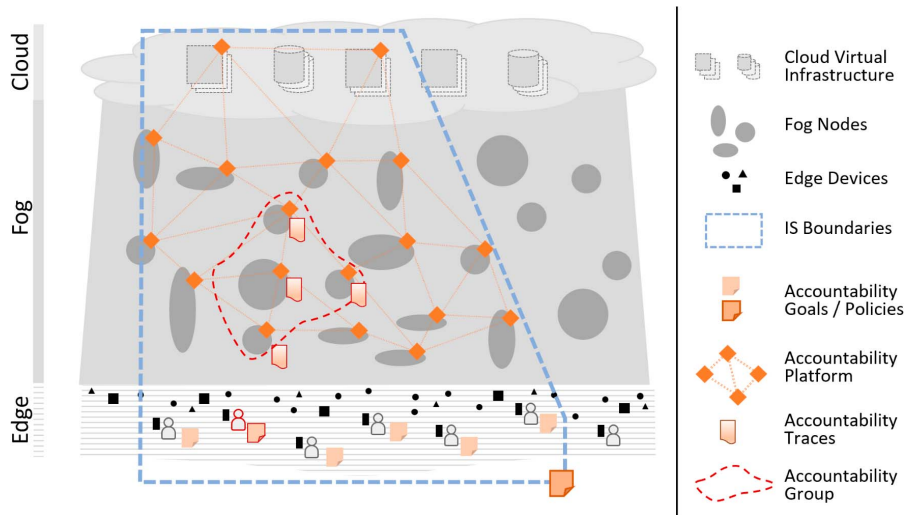


FIGURE 1. Abstract architecture of the FIDES perspective over the compute continuum.

an SLO in this context could represent the availability of an information system, explicitly defined as *the percentage of successfully served requests over the failed ones for a specific “check alive” endpoint, measured hourly*. Additionally, such SLOs are guaranteed to be bound to a specific range, ensuring that they meet a certain threshold, such as being above 99.9%.

In alignment with the specified outcomes, the FIDES requirements model is designed to express a comprehensive set of rules or policies. These rules serve to drive potential contingency plans that would be activated in case the expected outcomes are at risk. Notably, these contingency plans often entail the collaboration of multiple actors. To facilitate this collaborative effort, the concept of “accountability groups” is introduced. These groups consist of dynamic sets of participants that interact within the pre-existing choreography, allowing for coordinated and efficient responses to safeguard accountability and achieve desired outcomes.

For example, in the large AI models scenario discussed in the “Artificial Intelligence Revolution” section, a possible use case for an accountability requirements model is to ensure high accuracy for a specific question within a system where multiple AI systems are involved. For instance, in the context of a legal query or summarizing scientific articles, the accountability model could aim to achieve precise and reliable answers. When there is divergence among the AI systems’ responses, the accountability model would trigger the formation of an accountability group (see Figure 1). This group would expand the list of AI systems involved, seeking to verify a consensus on the

results and ultimately enhance the overall accountability and reliability of the system’s output.

In the context of the custom health systems discussed in the “Custom Health Care” section, accountability requirements models hold the potential to deliver personalized and health-conscious recommendations to users. For instance, when a user requests the system to suggest a place to eat, the model would take into account the user’s health preferences and requirements, ensuring that the recommendations align with their health goals and dietary restrictions. Additionally, the accountability model would verify the reliability and authority of the sources of nutrition information used by the system, steering clear of untrustworthy personal blogs or sources lacking proper references. By integrating such accountability measures, the custom health system can provide trustworthy and tailored recommendations, promoting users’ well-being and fostering responsible health guidance.

A service ecosystem scenario, such as the one described in the “Service Chains” section, and which focuses on the average response time from a specific integrated API, can also benefit from an accountability requirements model. In this case, the model aims to ensure that the average response time meets predefined performance standards. In the event of API performance degradation, the accountability model would trigger a contingency plan involving the creation of an accountability group tasked with scaling up the nodes that handle the API in question. By allocating additional resources to the API, the group aims to reduce the load and guarantee the desired response time,

thereby maintaining the chain's overall performance and reliability to meet user expectations.

On the other hand, it is essential to retain specific information to analyze and comprehend accountability status for users, allowing them to trace the accountability information involved in the system's operation. To address this need, the FIDES model incorporates a mechanism to define the accountability traces that are associated with each system's operation. For example, in the scenarios we have discussed, such traceability information may encompass a list of AI systems participating in a particular request and the verification of their responses. It could also include the sources of information used in health recommendations as well as the time stamps of requests and responses observed by the participants of a service ecosystem.

FIDES Accountability Platform

The FIDES accountability platform plays a pivotal role in managing and enforcing the specified accountability policies and models. This management entails a transparent and decentralized mechanism for creating, evolving, and distributing models among the various participants with access to the platform. Additionally, the platform takes responsibility for computing accountability metrics and evaluating adequacy of the outcomes. It should also provide a system for deploying and enforcing contingency rules or policies, collaborating with an accountability group management that guides the contingency choreography.

As an emerging trend in the industry, peer-to-peer meta-protocols or blockchain approaches (see the next section) are viable options for materializing these requirements. These technologies offer decentralized and transparent frameworks, which align well with the objectives of the FIDES accountability platform. By leveraging such innovative solutions, the platform can effectively ensure the integrity and reliability of the accountability mechanisms while promoting trust and collaboration among all participants involved.

APPLYING FIDES—THE FALCON CASE

To bring the FIDES framework into reality, we successfully developed an initial prototype⁵ inspired by a practical scenario centered around a public administration delivering citizen-oriented services. These services are facilitated through a service chain, consisting of diverse administration departments forming a federated infrastructure of interconnected information systems. Given this context, the significance of managing accountability becomes paramount as public administrations

strive to uphold transparency and efficiency as typical and essential requirements.

The FIDES prototype extends the Falcon framework,⁶ which monitors microservices infrastructures to ensure compliance with pre-established SLAs. Falcon collects metrics from the service infrastructure and computes the corresponding guarantees defined in the SLAs, providing multiple views of its fulfillment. The data and analyses obtained by Falcon play a crucial role in achieving the desired accountability, especially when addressing SLA violations. This makes it an ideal scenario for validating FIDES.

The developed FIDES prototype effectively tackles these challenges, offering a promising solution to enhance accountability within the federated infrastructure. It promotes responsible and transparent practices while ensuring efficient service delivery to citizens. By applying the FIDES approach to the Falcon framework, the management of accountability within the federated infrastructure experiences significant improvement, fostering trust and reliability in public administration services.

To realize the FIDES accountability requirements model (see the "FIDES Accountability Requirements Model" section), we leverage the SLA specification language from Falcon⁷ to define the various metrics and guarantees required for monitoring and calculations. The FIDES prototype extends this specification language, enabling the selection of specific metrics and guarantees from the comprehensive SLA covering the entire federated infrastructure. This selective approach ensures that only the data and analyses relevant to stakeholders' accountability needs are considered within the FIDES prototype.

The chosen metrics and guarantees, now included in the augmented SLAs, are stored in Falcon's registry component. Although Falcon's registry handles SLA management, the responsibility for accountable SLAs is seamlessly delegated to the equivalent registry component within the FIDES prototype. This seamless integration enables efficient management of accountable SLAs, promoting transparent accountability practices while focusing solely on the essential data and analyses required for stakeholders' needs.

During development of the first FIDES prototype, we opted for an instantiation of a permissioned blockchain as the foundation of the FIDES accountability platform (see the "FIDES Accountability Platform" section), which harnesses its capabilities to support metrics computation and associated analyses. Specifically, we established a blockchain network using Hyperledger Fabric,⁸ enabling the execution of required processes as smart contracts (referred to as *chaincode* in

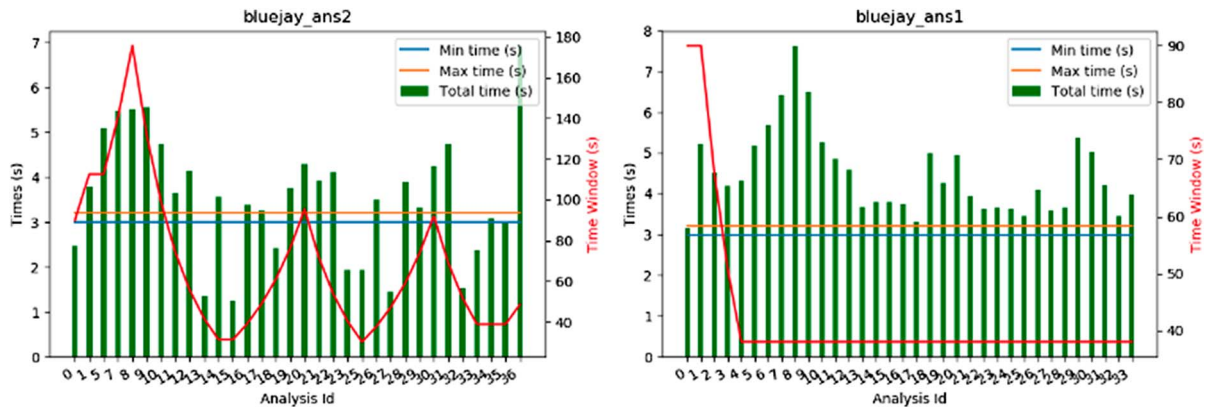


FIGURE 2. Validating FIDES execution over Falcon. min: minimum; max: maximum.

Hyperledger Fabric terminology). To seamlessly integrate the Falcon framework with the blockchain, we implemented a bridge using an accountable registry version of the original Falcon component.⁹

Through this approach, every metric collected by Falcon and each analysis performed to compute accountable guarantees, as chosen in the SLA, are registered and executed within the blockchain network. This integration ensures the desired levels of trust, transparency, and accountability for all the events related to the monitoring and enforcement of SLAs in the federated cloud infrastructure. The use of a blockchain-based solution strengthens the reliability and integrity of the accountability platform, further fostering confidence and transparency in the entire system.

Furthermore, the FIDES prototype utilizes an elastic smart contracts framework,¹⁰ which offers dynamic adaptation of process execution based on the complexity of the analyses and the number of metrics to be computed. This adaptive feature optimizes performance and efficiency of the system, ensuring seamless scalability to handle various scenarios and demands. The combination of a permissioned blockchain, the accountable registry, and the elastic smart contracts framework solidifies the FIDES prototype's robustness and reliability, making it an ideal solution for achieving accountability in federated cloud infrastructures.

Furthermore, we conducted a validation of our prototype by monitoring a microservice-based infrastructure known as *Bluejay*.¹¹ Bluejay is utilized to assess coordinated teams and integrated with multiple external services and tools. Consequently, performance of this solution may be affected depending on the number of tools and teams to coordinate. In this scenario, we established an SLA that guarantees a specific level of response time for Bluejay's operation by carefully

selecting both the associated metric and the guarantee defining the acceptable threshold as accountable elements for FIDES to consider. We then deployed the complete infrastructure using FIDES and conducted a series of experiments to demonstrate the feasibility and performance of our approach.

In Figure 2, we present the results of one such experiment where we executed the FIDES platform involving multiple instances within a federated deployment of Bluejay. The green bars illustrate execution time of the analyses for computing the guarantees associated with the two SLAs linked to Bluejay's operation. We set minimum and maximum thresholds for this time, represented by blue and orange lines in the figure. As the elastic smart contract detects that these thresholds have been surpassed, the time window (represented by red curves) adjusts to include fewer or more metric values for the subsequent guarantee computation accordingly. Thus, we observe how the FIDES performance varies throughout the execution, depending on the elasticity measures in place, thereby validating the feasibility of the platform to achieve federated accountability.

CONCLUSION

Our intention in this article was to delve into the risks associated with today's cloud-centric approach to information systems amid the profound societal transformations taking place. These risks encompass issues like privacy concerns, potential malicious usage, opaque biases, and reliability challenges, among others. To address these pressing issues, we proposed a paradigm shift, emphasizing accountability as a paramount aspect in the design and operation of information systems. Embracing a federated approach through the compute continuum paradigm, we presented the abstract framework of FIDES, which serves as a high-level blueprint

for designing and implementing federated accountability management in the compute continuum.

Grounding these abstract concepts, we developed a first prototype that materializes the FIDES framework by utilizing blockchain and SLA management technologies. This prototype serves as a federated accountability management platform for distributed infrastructure monitoring, providing a tangible demonstration of the potential of FIDES in real-world scenarios. By emphasizing accountability as a first-class citizen in information systems and embracing the compute continuum approach, we pave the way for more secure, transparent, and efficient management of distributed infrastructures.

In this context, it is important to acknowledge that our prototype represents a preliminary step, showcasing some FIDES elements while also presenting research opportunities for further exploration. Over the past few years, notable approaches have emerged that could address these challenges and complement FIDES. For instance, federated learning¹² offers a promising solution by combining models and datasets involved in AI generation while preserving user privacy, which could be applied to address accountability model distribution and evolution.

Moreover, OpenAI ChatGPT plug-ins¹³ provide customization and control over ChatGPT's behavior, enabling the gathering of accountability evidence and incorporating accountability requirements. Additionally, osmotic computing approaches¹⁴ align seamlessly with the compute continuum, creating dynamic computing ecosystems that collaborate to achieve specific objectives. These techniques could be leveraged to develop accountability groups and execute contingency plans when accountability outcomes are at risk. By embracing and integrating these novel approaches into future releases of FIDES instantiations, we can enhance its effectiveness and applicability, unlocking new possibilities for accountable and responsible information systems in the compute continuum.

ACKNOWLEDGMENT

This work was supported in part by Grant PID2021-126227NB-C21 and Grant PID2021-126227NB-C22, funded by MCIN/AEI/10.13039/501100011033/FEDER and European Union and Grant TED2021-131023B-C21 and Grant TED2021-131023B-C22, and funded by MCIN/AEI/10.13039/501100011033 and the European Union's NextGenerationEU/PRTR.

REFERENCES

1. D. Balouek-Thomert, E. G. Renart, A. R. Zamani, A. Simonet, and M. Parashar, "Towards a computing continuum: Enabling edge-to-cloud integration for data-driven workflows," *Int. J. High Perform. Comput. Appl.*, vol. 33, no. 6, pp. 1159–1174, Sep. 2019, doi: [10.1177/1094342019877383](https://doi.org/10.1177/1094342019877383). [Online]. Available: <https://journals.sagepub.com/doi/10.1177/1094342019877383>
2. V. Casamajor, P. K. Donta, A. Morichetta, I. Murturi, and S. Dustdar, "Edge intelligence—Research opportunities for distributed computing continuum systems," *IEEE Internet Comput.*, vol. 27, no. 4, pp. 53–74, Jul./Aug. 2023, doi: [10.1109/MIC.2023.3284693](https://doi.org/10.1109/MIC.2023.3284693). [Online]. Available: <https://ieeexplore.ieee.org/document/10184183>
3. "EU AI act: First regulation on artificial intelligence," *News Eur. Parliament*, Jun. 2023. Accessed: Jul. 27, 2023. [Online]. Available: <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
4. "Blueprint for an AI bill of rights: Making automated systems work for the American people," The White House, Washington, DC, USA, 2022. Accessed: Jul. 27, 2023. [Online]. Available: <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
5. "FIDES prototype." Zenodo. Accessed: Jul. 31, 2023. [Online]. Available: <https://doi.org/10.5281/zenodo.8197450>
6. "Falcon framework." GitHub. Accessed: Jul. 29, 2023. [Online]. Available: <https://github.com/governify/falcon-infrastructure>
7. "iAgree SLA specification language." Governify. Accessed: Jul. 31, 2023. [Online]. Available: https://www.governify.io/reference-guides/iAgree-5_2
8. "Hyperledger fabric." Hyperledger Foundation. Accessed: Jul. 29, 2023. [Online]. Available: <https://www.hyperledger.org/use/fabric>
9. "Falcon extension for smart contracts." Zenodo. Accessed: Jul. 31, 2023. [Online]. Available: <https://doi.org/10.5281/zenodo.8197446>
10. S. Dustdar, P. Fernandez, J. M. García, and A. Ruiz-Cortés, "Elastic smart contracts in blockchains," *IEEE/CAA J. Autom. Sin.*, vol. 8, no. 12, pp. 1901–1912, Dec. 2021, doi: [10.1109/JAS.2021.1004222](https://doi.org/10.1109/JAS.2021.1004222).
11. C. García et al., "Bluejay: A cross-tooling audit framework for agile software teams," in *Proc. IEEE/ACM 43rd Int. Conf. Softw. Eng., Softw. Eng. Educ. Training (ICSE-SEET)*, Madrid, Spain, 2021, pp. 283–288, doi: [10.1109/ICSE-SEET52601.2021.00038](https://doi.org/10.1109/ICSE-SEET52601.2021.00038).
12. B. McMahan and D. Ramage. "Federated learning: Collaborative machine learning without centralized training data." Google Research. Accessed: Jul. 31, 2023. [Online]. Available: <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>
13. "ChatGPT plugins." OpenAI. Accessed: Jul. 31, 2023. [Online]. Available: <https://openai.com/blog/chatgpt-plugins>

14. L. Carnevale, A. Celesti, A. Galletta, S. Dustdar, and M. Villari, "Osmotic computing as a distributed multi-agent system: The body area network scenario," *Internet Things*, vol. 5, pp. 130–139, Mar. 2019, doi: [10.1016/j.iot.2019.01.001](https://doi.org/10.1016/j.iot.2019.01.001).

AMADOR DURÁN is an associate professor at the University of Seville, 41004, Seville, Spain, and a member of the Ingeniería del Software Aplicada (Applied Software Engineering) Research Group and the Smart Computer Systems Research and Engineering. Lab. Contact him at amador@us.es.

PABLO FERNÁNDEZ is an associate professor at the University of Seville, 41004, Seville, Spain, and a member of the Ingeniería del Software Aplicada (Applied Software Engineering)

Research Group and the Smart Computer Systems Research and Engineering. Lab. Contact him at pablofm@us.es.

JOSÉ MARÍA GARCÍA is an associate professor at the University of Seville, 41004, Seville, Spain, and a member of the Ingeniería del Software Aplicada (Applied Software Engineering) Research Group and the Smart Computer Systems Research and Engineering. Lab. Contact him at josemgarcia@us.es.

SCHAHRAM DUSTDAR is a full professor of computer science and heads the Research Division of Distributed Systems at Vienna University of Technology, 1040, Vienna, 1040, Austria. Contact him at dustdar@dsg.tuwien.ac.at.

Computing in Science & Engineering

The computational and data-centric problems faced by scientists and engineers transcend disciplines. There is a need to share knowledge of algorithms, software, and architectures, and to transmit lessons-learned to a broad scientific audience. *Computing in Science & Engineering (CiSE)* is a cross-disciplinary, international publication that meets this need by presenting contributions of high interest and educational value from a variety of fields, including physics, biology, chemistry, and astronomy. *CiSE* emphasizes innovative applications in cutting-edge techniques. *CiSE* publishes peer-reviewed research articles, as well as departments spanning news and analyses, topical reviews, tutorials, case studies, and more.

Read *CiSE* today! www.computer.org/cise



IEEE
COMPUTER
SOCIETY

