# Towards Local Cloud Infrastructure in Developing Countries as a Response to Data Localization Regulations

Vladimir Inđić, Marija Kovačević, Miloš Simić, Goran Sladić
Faculty of Technical Sciences, University of Novi Sad, Novi Sad, Serbia
{vladaindjic, marija.kovacevic, milos.simic, sladicg}@uns.ac.rs

***Abstract*—Successful application of data localization laws is possible only if the country has an IT infrastructure mature enough to store all sensitive data, which is usually not the case in developing countries. This paper proposes a model for building local cloud infrastructure in developing countries compliant with data localization laws. The infrastructure will be developed gradually by adding location-aware commodity hardware nodes placed within the country's borders over a more extended period. Each node is considered a volunteer for a period defined by the participation policy. This kind of policy describes the formation of a cluster of geographically close nodes that might all be disposed of and replaced after the defined period expires. The government authorities take responsibility for motivating national universities and firms that do business inside the country's borders to take the key role in provisioning computing and storage nodes for a more extended period specified by the participation contracts and policies. In return, authorities guarantee a set of benefits such as tax reductions and funds for R&D. The authorities should make laws and regulations that follow the gradual development of the local cloud infrastructure by introducing soft data localization first. Hard data localization will replace soft data localization at the moment of infrastructure maturity.**

***Keywords*—local cloud infrastructure, developing countries, data localization regulations, data localization laws, location-aware nodes, participation policies, participation contracts**

## I. INTRODUCTION

The decreasing trend of prices for "pay-as-you-go" services, combined with easy on-demand workload scaling, made cloud computing extremely popular among SMEs and multinational companies. In the past decade, developed country governments raised concerns about whether a potential conflict with another country's government may lead to losing access to the data stored in cloud data centers situated on the territory under the jurisdiction of that country's government. Developed country governments made a set of laws and regulations about data localization that forbid sensitive data from leaving the country's borders to prevent potential misuse of the data by governments of countries in conflict. Those laws are data localization or data sovereignty laws.

A mature IT infrastructure that can store all sensitive data is required to apply the data localization laws successfully. Developing countries usually fail to satisfy this requirement. It would be unreasonable for a developing country's government to pass the data localization legislation before the moment when a significant number of local firms can fulfill the laws' requirements. Instead of expecting that local firms are ready to allocate funds for assembling private IT infrastructures that are laws compliant, the government authorities should put an effort to make a plan for the development of the unique national-scale IT infrastructure that local firms could use for hosting their services and data at reasonable costs.

Thus, we propose a local cloud infrastructure development model considering three important aspects. The first aspect represents a technical design of the infrastructure given in the form of the extended model for the dynamic organization of geo-distributed nodes into disposable clusters initially described by Simic et al. [1]. The second aspect focuses on the provisioning strategy involving collaboration among government authorities, IT companies, and national universities to provide volunteer nodes and people to work on their administration. The last important aspect emphasizes the importance of coordinating laws restrictiveness towards data flows with the potential of the local cloud infrastructure to support required restrictiveness successfully.

### A. *The Document Roadmap.*

The following section describes the current research on data localization laws. Section 3 gives the technical design of the local cloud infrastructure development model compliant with data localization laws. Section 4 presents the node provisioning strategy as a prerequisite for practical infrastructure development. Section 5 describes a proper way to coordinate data localization laws implementation with local cloud infrastructure development. The last section highlights the drawbacks of data localization laws and advocates for the free flow of depersonalized non-sensitive data.

## II. BACKGROUND

This paper briefly elaborates on work done in data protection and localization laws and the IT infrastructure developed for their successful application. In 2018, European Union implemented General Data Protection Regulation to increase the safety of sensitive personal data of EU citizens. It is done by restricting the flow of sensitive data outside of the EU [2]. Since then, GDPR has been evolving into the world's most significant de facto data localization framework [3], which led to emerging of the GDPR compliant distributed cloud

infrastructure dispersed across the EU territory known as GAIA-X [4].

GDPR served as a role model for designing the Personal Information Protection Law [5], which has been recently passed by China government, infamous for its orientation toward data restrictiveness. Russia is another data-restrictive country that significantly contributed to the data localization laws field. In 2014, as a response to the international sanctions that targeted Crimea-based services that forced Visa and Mastercard to end services there, Russia required payments data localization as part of an initiative to create a Russian payment system called MIR [3].

The number of data localization laws is growing in developing countries. As the developing world leaders, India and Turkey proposed a significant number of data localization policies, 12 and 7, respectively. In 2012, India enacted the "National Data Sharing and Accessibility Policy," requiring government data to be stored in local data centers [3]. The Reserve Bank of India has adopted rules requiring all payments to stay within the country's borders in 2018 [6].

Two years later, Turkey adopted the "Law on Amendment of the Law on the Regulation of Publications on the Internet and Suppression of Crimes Committed by Means of such Publications" which includes data localization and grants the government more power to regulate content on social media [3]. Namely, social network providers with more than 1 million users must keep the personal data of Turkish citizens in the country [3]. Further, Serbia adopted the Law on Protection of Personal Data in 2018 [7]. Serbia is currently in the process of EU integration, which is why this law is based on the GDPR. Thus, the provisions of the law comply with the requirements of the GDPR.

Azmeh and Foster [8] pointed out two important benefits that developing countries can derive from data localization policy. The first one represents a higher foreign direct investment in digital infrastructure. The second benefit is the impact of server localization on the local infrastructure development and the presence of skilled professionals that can lead to the growth of the local IT industry.

India shows interest in developing infrastructure compliant with data localization laws to support their successful application. The committee of Experts under the Chairmanship of Justice B.N. submitted a report [9] recommending the development of the local cloud infrastructure that enables small local firms to store sensitive data to comply with data localization and protection laws. We go one step further by proposing a local cloud infrastructure development model that gradually introduces volunteer nodes dynamically organized into disposable clusters. We also provide guidelines to the government authorities about the node provisioning strategies and how to coordinate the process of the laws' implementation with the infrastructure development.

## III. TECHNICAL SOLUTION MODEL

An integral part of traditional cloud infrastructure represents a data center equipped with high-end computing, network, and storage hardware [10]. Building these data centers might be a vast starting step for the government authorities of developing countries because of the initially high expenses [11]. Instead, it is more likely to develop the infrastructure gradually by adding small clusters or even single nodes that might not be placed within the same facility but rather dispersed over the territory of one or multiple cities of a country. To minimize initial investments, nodes can be built by using commodity hardware. Since the developing countries might not always allocate enough system administrators to work on the infrastructure maintenance, a group of nodes may go down or be purposefully excluded from the infrastructure at some critical moments. Thus, all nodes should be considered volunteer nodes that can join the infrastructure at some point in time and serve the users' requests until the moment of graceful exclusion from the infrastructure. Both joint and excluding moments must be defined in front by related policies that each of the nodes, as well as system administrators, should respect.

The gradual nature of the development process may introduce heterogeneity in the hardware used to build the infrastructure. Integrating the heterogeneous nodes dispersed across multiple location sites is not a trivial job. To the best of our knowledge, current open source cloud platforms such as *OpenStack* require significant extensions to support decentralized data-centers integration [12]. However, the model for the dynamic organization of geo-distributed nodes into micro data centers used to build distributed cloud architectures proposed by Simic et al. [1] might be applicable. For brevity, we refer to this model as the µC. Any set of geographically close nodes connected to the network and employed to do the same job may be considered a physical cluster. One or more clusters placed over an arbitrary geographic region form a logical region, while multiple regions form the topology. Regions can dynamically reduce the size by adding new clusters or disposing of the current ones to respond to the change in the size of the system's workload. If the entire cluster goes down, the region can failover the workload to the cluster of the same region to preserve the system's availability. A similar stands for the topology. To form a new cluster, a user defines a query by giving a set of nodes' attributes called labels used to identify requested nodes. The system responds to the query by forming a new cluster if all of the identified nodes are idling in the pool of free nodes. This communication flow, depicted in Fig. 1, is defined as the cluster formation protocol whose consistency is formally proven.

We introduce two important labels that each of the

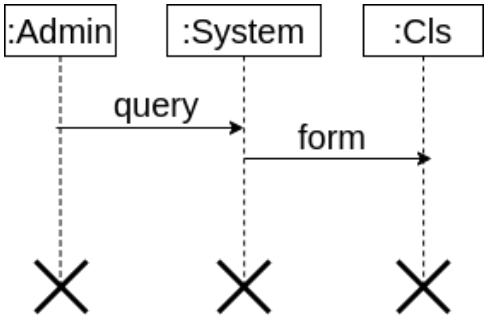


Figure 1. The communication diagram corresponds to the simplified cluster formation protocol of the µC model that creates the *Cls* cluster [13].

nodes must have. The first is a geo-location of a node that must represent a location site within the country's borders. The second label represents the participation policy specifying the time interval during which the node volunteers as a part of the infrastructure. Clusters should be formed of nodes that are geographically close to each other and whose end of the volunteering time interval matches. The last property of the nodes in the cluster can automate the process of disposing of the whole cluster from the infrastructure instead of excluding each node separately when the volunteering date expires. If a user of the local cloud infrastructure still needs to use services that run on the disposed cluster, the disposal process must be done gracefully. Closely before the moment of actually disposing of a cluster A, while this cluster might still be active, all services running on A must be slowly migrated to another cluster B of the infrastructure topology, followed by the redirection of the user's requests to the cluster B. Thus, the graceful disposal process is considered a failover of the cluster's workload that must not harm the user data while providing the highest possible level of system availability.

Besides the operation responsible for creating a new cluster, the μC model supports operations for editing or deleting a cluster. All three operations are mutated operations. However, an operation that could move a workload from an ending cluster A to another already running cluster B is not supported. Thus, we propose introducing the MOVE operation that could enable a graceful shutdown [14]. The MOVE operation allows the user to provide a cluster identifier whose workload will be moved to another similar idling cluster. The cluster with the specified id will be deleted from the architecture eventually.

Suppose the service running on a cluster A before gracefully shutting down is stateless and does not require data transfer. In that case, the MOVE operation is simple. It involves starting an instance of the same service on a cluster B and rerouting the network traffic from A to B. On the contrary, the MOVE operation needs to transfer all data from a cluster A to B. Before the transfer happens, cluster A is declared as temporarily unreachable not to acquire more data nor accept any future requests. As in the previous case, the MOVE operation instantiates a service run on cluster B and reroutes all network traffic after the transfer finishes.

One can notice an interval during which service running on an ending cluster might be unresponsive to the end-user, which is unacceptable for critical systems. Hence, we suggest an extension to MOVE operation applicable to critical services. A critical service running on an arbitrary cluster can asynchronously back up the data to another cluster to increase service reachability. Thus, the MOVE operation can reroute the network transfer to a backup cluster with minimal unresponsive time. As one may notice, a close correspondence between critical services designers and system administrators maintaining cloud infrastructure is inevitable. As Simic et al. proposed [13], introducing a new policy specifying additional needs that a critical service depends upon represents a unified way for successful correspondence.

It is tedious for system administrators to keep track of clusters whose volunteer period expires and transfer their workload via MOVE operation. Thus, we propose a way to automate this process by introducing a new component

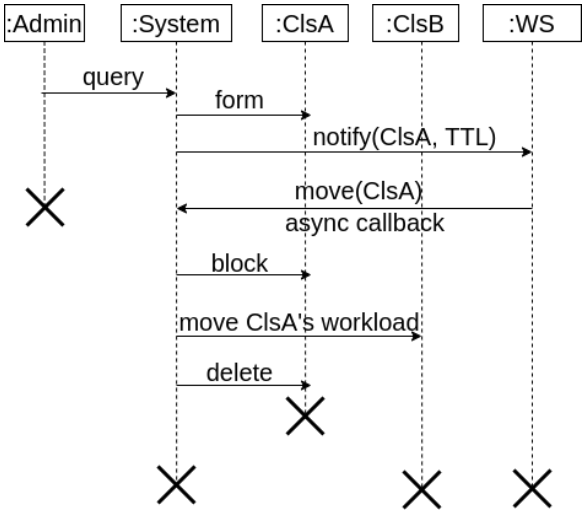


Figure 2. The cluster formation protocol, depicted in Figure 1, is extended with the Workload shifter component (*WS*). This component is responsible for asynchronously initiating MOVE operation that results in moving the cluster A's (*ClsA*) workload to another running cluster B (*ClsB*) and deleting the cluster A (*ClsA*).

in the system depicted in Fig. 2. The Workload Shifter component is an arbitrary key-value database that supports the time to live (TTL) mechanism [15], allowing one to specify a period after an entry stored in the database expires and gets automatically removed from the storage. Entry removal is directly followed by a trigger of a corresponding expiration callback. We extended the cluster formation protocol shown in Fig. 1 in the following way. Namely, after the system forms a new cluster, it sends the cluster id and the end of the cluster volunteering period to the Workload Shifter component. After the period expires, the component automatically removes the id from the database and triggers the corresponding expiration callback. We use the callback to send the MOVE operation to the system with the id of the cluster whose volunteering period expired. As a response, the system moves a workload from the ending cluster to another similar running cluster, followed by removing the cluster with the given id.

## IV. PROVISIONING STRATEGY

The developing country's government is responsible for strategically planning a gradual development of the infrastructure by focusing on providing and maintaining nodes in clusters. A special IT agency dedicated to technical development should be formed and ready to employ people to work on node provisioning and administration. The budget of a developing country usually does not have sufficient funds to support the whole infrastructure development by the agency alone. Instead, an initial set of nodes should be provided for the infrastructure by the agency that receives a grant from the country's budget. At the same time, the government should try to involve other interested parties in infrastructure development.

Lower electrical power prices in developing countries were one of the key arguments for building data centers by multinational IT companies in those countries. Likewise, lower salaries lead to an increasing number of local people in the IT industry working on outsourcing projects for foreign companies. This implied the growth of the local IT companies in the developing countries in the

last decade. All companies that do business inside the country's borders that involve data declared as sensitive by the data localization laws, whether their facilities are situated inside or outside of its borders, must respect the data localization laws. Mature companies might have enough funds to form their own on-site data localization laws compliant solutions, which is not the case for the most SMEs and start-up companies.

The government should involve matured IT companies to support local cloud infrastructure development. It should design participation contracts that specify benefits for the firms employed in provisioning some sets of volunteer nodes during the contracted period. A significant benefit for the companies is tax reductions. After the expiration of a participation contract, if a company decides not to be involved in the development process anymore, all of the provisioned cluster nodes are automatically disposed of and revoked from the infrastructure owing to the participation policies installed on the nodes.

SMEs and other non-IT firms might not have financial interests in building private on-site solutions. Instead, they can rent the local cloud infrastructure nodes and pay for the used services. However, start-up companies might not have enough funds to allocate for renting nodes of the local cloud. A start-up company might receive government benefits in the form of specific renting contracts that define a period during which some set of nodes is rented to the start-up company for free. After the due date, when the start-up company matures enough, it is obliged to either pay a predefined amount of money to cover the renting expenses or provide volunteer nodes to the local cloud instead. Besides strategies for node provisioning, it is important to consider providing enough engineers to work on infrastructure administration. Companies offering volunteer nodes might allocate some of their system administrators' working hours for node maintenance. Universities could take a key role in providing enough technical resources for cloud administration. Instead of teaching, younger employees might work on cluster maintenance. Involving graduate students in infrastructure maintenance as part of the DevOps or similar courses can improve the quality of studies because students could gain more practical experience during the lessons.

## V. Coordinating Data Localization Laws and Infrastructure Development

We can recognize two types of data localization: soft and hard [16]. The first one involves "mirroring", meaning that copies of a certain kind of data must be stored within a specific area but do not restrict that data's copying or transmission elsewhere [16]. The second type mandates that specific data, including critical infrastructure data and sensitive personal data, must be stored within certain borders and not allowed to be transmitted across borders [16]. We advise government authorities to gradually introduce a set of data localization laws and coordinate their applications with the potential of the local cloud infrastructure. Passing laws about soft data localization would be a reasonable starting step since the demands of this kind of localization should be well supported by the infrastructure in its early stages. After the infrastructure matures enough, more restrictive data localization laws might be considered, leading to the eventual introduction of hard data localization.

## VI. Conclusion

The number of data localization laws drastically increased in both developed and developing world countries in the past decade. Successful application of data localization laws is possible only if a country has an IT infrastructure mature enough to store all sensitive data, which is usually not the case in developing countries. This paper proposes a model for building local cloud infrastructure in developing countries compliant with data localization laws. Proposed infrastructure development should go gradually by adding location-aware commodity hardware nodes within the country's borders over a more extended period. We describe this model by considering three important aspects. The first aspect represents a technical design of the infrastructure given in the form of the extended µC model for the dynamic organization of geo-distributed nodes into disposable clusters initially described by Simic et al. [1]. The second aspect focuses on the provisioning strategy involving collaboration among government authorities, IT companies, and national universities to provide volunteer nodes and people to work on their maintenance. The last important aspect emphasizes the significance of coordinating laws restrictiveness towards data flows with the potential of the local cloud infrastructure to support required restrictiveness successfully.

The µC model we extended in this paper is in its early practical stages [1, 13], and it is still not ready for use in real-world applications, thus making our work purely theoretical. However, we provide a solid theoretical foundation for future practical implementation of local cloud infrastructure in developing countries with data localization laws compliance.

We want to point out that data localization introduces some drawbacks. One can misuse it and consider it as a replacement for data security, thus introducing vulnerabilities in local cloud infrastructure that could be exploited. Hence, we emphasize that all good data security practices have to be followed while implementing data localization laws compliant infrastructure. Storing data within a country's border might constrain data sharding on the country's area. Even unlikely, a disaster can hit the area of the data residence, thus causing persistent data loss. Also, data localization constrains the free data flow, which is essential for science development. Therefore, we advocate that at least depersonalized data should be exported outside a country to be stored in the global cloud and used in scientific research, as Simic et al. proposed [13]. The µC model allows this kind of export [13], meaning that the infrastructure model we describe easily integrates with the global cloud.

## References


[1] M. Simić, I. Prokić, J. Dedeić, G. Sladić, B. Milosavljević, Towards edge computing as a service: Dynamic formation of the micro data-centers, IEEE Access 9 (2021) 114468–114484.

[2] Data residency laws by country: an overview, https://incountry.com/blog/data-residency-laws-by-country-overview/, Accessed: 2022-01-13.

[3] N. Cory, L. Dascoli, How Barriers to Cross-Border Data Flows Are Spreading Globally, What They Cost, and How to Address

Them, Technical Report, Information Technology and Innovation Foundation, 2021.

[4] GAIA-X, https://www.data-infrastructure.eu/GAIAX/Navigation/EN/Home/home.html, Accessed: 2022-01-13.

[5] The great wall of data privacy: China passes comprehensive data privacy law, https://www.jdsupra.com/legalnews/the-great-wall-of-data-privacy-china-8338636/, Accessed: 2022-01-13.

[6] Statement on developmental and regulatory policies, https://rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=43574, Accessed: 2022-05-30

[7] Serbia - data protection overview, https://www.dataguidance.com/notes/serbia-data-protection-overview, Accessed: 2022-05-30.

[8] S. Azmeh, C. Foster, The TPP and the digital trade agenda: Digital industrial policy and Silicon Valley's influence on new trade agreements, Technical Report, Working Paper Series, 2016.

[9] A free and fair digital economy, https://www.meity.gov.in/writereaddata/files/Data_Protection_Committee_Report.pdf, Accessed: 2022-01-13.

[10] F. R. de Souza, C. C. Miers, A. Fiorese, M. D. de Assunção, G. P. Koslovski, Qvia-sdn: Towards qos-aware virtual infrastructure allocation on sdn-based clouds, Journal of Grid Computing 17 (2019) 447–472.

[11] M. Satyanarayanan, P. Bahl, R. Caceres, N. Davies, The case for vm-based cloudlets in mobile computing, IEEE pervasive Computing 8 (2009) 14–23.

[12] A. Lebre, J. Pastor, A. Simonet, F. Desprez, Revising openstack to operate fog/edge computing infrastructures, in: 2017 IEEE international conference on cloud engineering (IC2E), IEEE, 2017, pp. 138–148.

[13] M. Simić, G. Sladić, M. Zarić, B. Markoski, Infrastructure as software in micro clouds at the edge, Sensors 21 (2021) 7001.

[14] F. P. Tso, S. Jouet, D. P. Pezaros, Network and server resource management strategies for data centre infrastructures: A survey, Computer Networks 106 (2016) 209–225.

[15] E. Cohen, E. Halperin, H. Kaplan, Performance aspects of distributed caches using ttl-based consistency, Theoretical computer science 331 (2005) 73–96.

[16] Data localization: No panacea for cloud computing issues, https://lawschoolpolicyreview.com/2021/08/18/data-localization-no-panacea-for-cloud-computing-issues/, Accessed: 2022-01-13.