

LDOP-Edge Architecture for Distributed Data Processing and Sharing in E-Government

DOI: <https://doi.org/10.54654/isj.v2i28.6418>

Pham Van Dinh, Ha Mai Huyen Trang*, Pham Thi Son, Vu Ngoc Duong

Abstract—E-government systems use growing volumes of data across agencies and administrative levels. In a fully centralized model, each local request traverses a wide-area network (WAN), which may increase latency, consume bandwidth, and interrupt point-of-service operations during outages. This design paper proposes LDOP-Edge, a governed edge layer integrated with the National Data Platform (NDOP), secure connection points (AGN), local data platforms (LDOP), and authoritative source systems. LDOP distributes machine-enforceable policies for cache eligibility, time-to-live (TTL), permitted degraded-mode operations, access control, retention, synchronization, and evidence logging. Edge Nodes store only approved minimum data and never become authoritative sources. Cached records carry source, version, update time, validity, integrity protection, and verification status. Offline transactions carry event and idempotency identifiers and remain provisional until central validation. The paper contributes a public-sector architecture mapping, a decision matrix for policy-controlled caching, a transaction state machine, a threat-control model, and a reproducible protocol for comparing centralized, simple-cache, and LDOP-Edge configurations. Because the current work is architectural and no measured prototype results are claimed, expected performance and continuity benefits are stated as hypotheses for controlled validation.

Tóm tắt— Các hệ thống chính phủ điện tử khai thác khối lượng dữ liệu ngày càng lớn giữa nhiều cơ quan và cấp hành chính. Trong mô hình tập

trung hoàn toàn, mọi yêu cầu tại đơn vị cơ sở phải truyền qua mạng diện rộng (WAN), có thể làm tăng độ trễ, tiêu tốn băng thông và gián đoạn nghiệp vụ khi mất kết nối. Bài báo đề xuất kiến trúc LDOP-Edge, trong đó LDOP quản lý dịch vụ dữ liệu, chính sách chia sẻ, kiểm soát truy cập và đồng bộ; các nút biên được triển khai gần người dùng để lưu đệm dữ liệu đã được phê duyệt, xử lý cục bộ và tạm tiếp nhận một số giao dịch khi mạng WAN gián đoạn. Nút biên không phải là nguồn dữ liệu có thẩm quyền và chỉ được lưu lượng dữ liệu tối thiểu theo chính sách. Mỗi bản ghi có thông tin về nguồn, phiên bản, thời điểm cập nhật, thời hạn hiệu lực, cơ chế bảo vệ tính toàn vẹn và trạng thái xác minh. Giao dịch ngoại tuyến có mã sự kiện, khóa chống xử lý trùng và chỉ có hiệu lực chính thức sau khi được hệ thống nguồn xác thực. Đóng góp của bài báo gồm ánh xạ kiến trúc dữ liệu khu vực công, ma trận quyết định lưu đệm và TTL, máy trạng thái giao dịch, mô hình đe dọa - kiểm soát và quy trình kiểm chứng có thể tái lập. Do nghiên cứu hiện tập trung vào thiết kế kiến trúc, các lợi ích về hiệu năng và tính liên tục được trình bày dưới dạng giả thuyết cần được kiểm chứng bằng thử nghiệm có kiểm soát.

Keywords— *E-government, LDOP, NDOP, AGN, edge computing, data sharing, distributed data, data governance.*

Từ khóa— *chính phủ điện tử, LDOP, NDOP, AGN, điện toán biên, chia sẻ dữ liệu, dữ liệu phân tán, quản trị dữ liệu.*

I. INTRODUCTION

A. Background and problem statement

E-government depends on timely, reliable, and well-governed data. Data is not only information stored in a database; it is also a public asset that must be protected and shared under clear rules. In a traditional centralized model, authoritative data is stored in central systems. A local unit sends each query or update

This manuscript was received on July 15, 2026. It was reviewed on August 12, 2026, revised on August 16, 2026 and accepted August 20, 2026.

*Corresponding author

over the network to the center. This model simplifies control of original data, but it creates three problems. First, more users and queries increase the load on central systems and WAN links. Second, network delay can slow operations that require an immediate response. Third, a WAN outage may prevent local units from accessing data and disrupt services at the point of use.

Edge computing offers a possible solution. It moves part of the storage and processing capacity closer to the place where data is generated or used. This can reduce response time and WAN traffic. However, a government Edge Node cannot freely copy or manage data as an independent repository. Government data is subject to legal, security, authorization, retention, and accountability requirements. The Edge must therefore operate as a controlled execution layer within the common e-government data architecture.

This paper proposes the LDOP–Edge architecture. LDOP is the data platform of a ministry, sector, or local authority. An Edge Node provides local processing close to users and data sources. NDOP is the National Data Platform, while AGN provides a secure connection between LDOP and NDOP. The architecture combines centralized data governance with controlled local processing. It does not replace authoritative source systems or create a separate local source of official data.

B. Related work and policy context

In e-government, data supports many agencies, administrative levels, and operational domains. Data interoperability is therefore more than a technical connection between systems. It also requires data governance, a clearly identified authoritative source, access control, logging, and traceability. In this paper, an authoritative source means the system that stores original data or data with official value.

In Vietnam, Decision No. 2439/QĐ-TTg introduced the National Data Architecture Framework. It identifies components such as NDOP, LDOP, AGN, master data, transaction data, aggregated data, analytical data, metadata,

and a shared data dictionary [1]. The 2024 Data Law, Decree No. 194/2025/ND-CP, and Decree No. 278/2025/ND-CP provide a further legal basis for managing, connecting, sharing, and using public-sector data [2 - 4]. Earlier, Decree No. 47/2020/ND-CP established the foundation for managing, connecting, and sharing digital data among public agencies [5].

Research on edge computing shows that moving processing, storage, and analytical capacity closer to data sources can reduce latency and traffic to central systems. It can also support applications that require rapid responses [6, 7]. The NIST fog computing model similarly organizes computing, storage, and networking resources into multiple layers. This structure suits systems with many data sources, including IoT devices, cameras, sensors, and monitoring equipment [8]. In the public sector, however, an Edge Node should not be treated as an independent data repository. Edge data should be limited to cached data, temporary data, or local processing results used within an approved scope. Each Edge record should include its source, update time, version, validity period, verification status, and permitted use. This approach captures the benefits of distributed processing without weakening the role of the authoritative data source [1, 2].

For data security, Attribute-Based Access Control (ABAC) grants permissions according to user, data, action, and access conditions [9]. Zero Trust Architecture assumes that no user, device, or service is automatically trusted, even within an internal network [10]. All Edge access must therefore be authenticated, authorized, logged, and controlled by policies managed through LDOP.

The OECD also states that a data-driven public sector should treat data as a strategic asset. Data can improve public services, support decision-making, and strengthen accountability [11]. This view is consistent with the LDOP–Edge architecture. The Edge provides faster and more flexible processing, while LDOP continues to coordinate policies, share data, control access, and maintain traceability.




Figure 1. Logical flow of national data sharing and coordination [1]

More recent work addresses safe local-first applications, secure replication, model checking of offline synchronization, and authenticity of edge caches [12 - 15]. Those works solve important technical subproblems, but they do not define how a governed Edge Node should fit into the NDOP-AGN-LDOP-source-system structure, how official authority must be preserved, or which evidence public agencies should retain for inspection and accountability.

Research published in this journal also provides useful evidence for distributed and privacy-aware processing. Ho et al. coordinate centrally managed tasks across a distributed computer network and return results to the management component [16]. Khanh et al. use federated learning for collaborative cloud workload forecasting without exposing local data, illustrating a privacy-preserving separation between local processing and global coordination [17]. However, neither study defines policy-controlled caching, authoritative public-sector records, degraded-

mode use, or reconciliation of provisional government transactions; these remain the focus of LDOP-Edge.

Previous studies explain the technical benefits of edge computing, while legal and policy documents define public-sector data-governance requirements. However, they do not provide an operational model that connects Edge Nodes to the NDOP-LDOP structure while preserving the authority of source systems. The specific contributions of this paper are: (1) a role model for NDOP, AGN, LDOP, source systems, and Edge Nodes; (2) a policy-controlled model for caching and limited degraded operation; (3) an auditable synchronization and reconciliation lifecycle that prevents the Edge from overwriting authoritative data; and (4) a proposed validation framework comparing centralized, simple-cache, and LDOP-Edge configurations. These are architecture and design contributions; their performance benefits remain to be verified through a controlled pilot.

TABLE I. COMPARISON OF REPRESENTATIVE ARCHITECTURE APPROACHES

Criterion	Centralized / generic Edge	LDOP–Edge
Authoritative source	Central store or application-specific	Source system remains authoritative
Caching policy	None or application-defined	LDOP-approved, minimum data and TTL
Offline operation	Unavailable or locally defined	Controlled degraded mode only
Sync/conflict	Not applicable or implementation-specific	Event ID, validation, reconciliation, review
Traceability	Central or platform logs	Source/version labels plus access and sync logs
Security control	General perimeter/device controls	Policy enforcement, mutual trust, signed events
Public-sector fit	Generic computing architecture	Mapped to NDOP–AGN–LDOP and legal authority

TABLE II. EXPLANATION OF KEY CONCEPTS

Concept	Plain-language meaning	Role in this paper
NDOP	National Data Platform	Connects, shares, and governs data nationwide
LDOP	Data platform of a ministry, sector, or local authority	Coordinates data sharing, distribution, and monitoring within its management scope
AGN	Secure connection point	Provides controlled connectivity between LDOP and NDOP
Edge Node	Processing node close to data users	Caches data, processes it locally, and accepts temporary transactions during outages
TTL	Validity period of cached data	Determines how long the Edge may use stored data
Authoritative source	System holding original or officially valid data	Ensures that Edge data does not replace original data

C. Research questions and requirements

The central question is not whether a distributed system should replace the centralized model. Instead, the architecture must answer four practical questions: Which data must remain centrally managed? Which data may be cached at the Edge? Which cases require only metadata or aggregated results to be sent to the center? What may a local unit do when the WAN connection is unavailable?

In practice, many operations do not require all raw data to be transferred to the center. Cameras and IoT devices, for example, may generate very large data streams, while the central system may need only alerts, relevant segments, metadata, or analytical results. By contrast, data that is critical, frequently updated, or legally significant should not be copied to the Edge for long-term storage. The Edge should retain only the minimum attributes or a verification token.

Based on these questions, the architecture must meet four requirements: preserve the authority of source systems; reduce latency for frequent queries; maintain selected local operations during WAN outages; and enforce security, logging, and accountability throughout the data life cycle.

TABLE III. LIMITATIONS OF THE CENTRALIZED MODEL AND EDGE-BASED RESPONSES

Limitation	Practical effect	How the Edge helps
High latency	Queries from local units travel over a distant WAN connection	Returns valid data directly from the Edge
High bandwidth use	Large volumes of raw data, video, images, and logs are sent to the center	Processes data locally and sends only alerts, metadata, or necessary results
Dependence on connectivity	Some operations stop when the WAN fails	Enables controlled degraded operation
Limited scalability	More users increase the load on the central system	Distributes part of the processing load

		to suitable nodes
Risk of uncontrolled copies	Unrestricted copying weakens data control	Stores only the minimum data approved by LDOP policy

II. PROPOSED LDOP-EDGE ARCHITECTURE

A. Design principles

The first principle is to preserve the role of the authoritative data source. An Edge Node may store copies, cached data, temporary transaction data, or analytical results. It must not become the final authority for original data.

The second principle is data minimization at the Edge. Only data that is necessary, relevant to the local area, within the user's authority, and subject to a clear retention period may be stored. This approach reduces the risk of data exposure and prevents the creation of uncontrolled distributed repositories.

The third principle is that every Edge record must carry an information label. At a minimum, the label includes the data source, update time, version, validity period, verification status, permitted scope of use, and deletion policy.

The fourth principle is that offline operation must be limited. If the WAN connection fails, the system may allow users to read data that remains

valid or temporarily submit selected transactions. These transactions enter a queue, receive unique event identifiers, and are protected against duplicate processing. A transaction is considered complete only after it has been synchronized and reconciled with the source system.

The fifth principle is that every access request must be authenticated, authorized, and logged. The logs should answer the following questions: who accessed the system, what data was accessed, when and why it was accessed, what the outcome was, and whether any data was transmitted externally.

B. Proposed LDOP-Edge architecture

The proposed architecture has four main architectural components. The first consists of NDOP and national-level data components. It supports nationwide data connectivity, sharing, aggregation, and governance. The second is AGN, which provides a secure connection between the national data platform and the data platforms of ministries, sectors, and local authorities. The third is LDOP. It coordinates data services, manages sharing policies, distributes data, records logs, and connects to source systems. The fourth is the Edge Node, deployed close to where data is generated or used. It caches approved data, processes data locally, and synchronizes data according to policy.




Figure 2. Overall LDOP-Edge architecture for data interoperability

Figure 2 translates these four components into the proposed logical architecture. It should be read from the national coordination layer through AGN and LDOP to source systems and Edge Nodes. In this architecture, an Edge Node does not operate independently. It receives policies from LDOP, including the approved data list, retention periods, authorized user groups, conditions for using data during outages, and procedures for synchronization after connectivity returns. The Edge improves performance and availability while remaining within the common governance framework.

Requests to access an Edge database must pass through a policy enforcement point governed by LDOP policy. As a result, the same authorization rules apply consistently at both the center and the Edge.

The framework-to-architecture mapping is as follows. NDOP supports national coordination; AGN provides the controlled interconnection; LDOP distributes policies and records access and synchronization evidence; source systems retain official authority and validate writes; and Edge Nodes enforce local policy without becoming authoritative repositories. Evidence retained for audit includes the policy version, source and record version, access decision, event ID, synchronization result, and any conflict-review decision.

III. DATA POLICIES AND PROCESSING FLOWS

A. Edge data policies

Not all data should be transferred to the Edge. This paper divides data into several classes, each with different storage, use, and synchronization policies. Table IV provides a reference classification for practical system design.

TABLE IV. DATA CLASSES AND THEIR PROCESSING AT THE EDGE

Data class	Processing at the Edge	Management requirements
Reference data and code lists	May be stored as versioned copies	Must be synchronized when a new version becomes available
Master data	Stores only minimum attributes, tokens, or short-term cached copies	Must identify the authoritative source and verification status
Local transaction data	May be temporarily queued during an outage	Must prevent duplicate processing and be reconciled after connectivity returns

High-volume IoT, camera, and log data	Processes data locally and sends only events, metadata, alerts, or necessary segments	Should avoid sending all raw data unless required
Analytical and AI data	Stores inference results, risk scores, and alerts	Must include the model version and result generation time
Metadata and logs	Must be recorded and synchronized with the monitoring system	Support tracing, inspection, and incident handling

TABLE V. DATA STATUSES DISPLAYED TO USERS

Status	Meaning	Permitted use
Verified online	Data has just been checked against the authoritative source	May be used in accordance with operational rules
Valid cached data	Data is stored at the Edge and its TTL has not expired	May be used for rapid queries; the update time must be displayed
Expired cached data	The validity period has ended	Must not support important decisions until reverified
Pending synchronization	Transaction was created during an outage	Remains provisional until reconciliation is complete
Conflict detected	Edge data differs from source data or duplicates another transaction	Requires handling through an operational review process

For master data, the Edge should store only the minimum information needed for rapid verification or a token that refers to the authoritative source. Local transaction data may be temporarily accepted during a network outage, but it must later be synchronized and reconciled. For cameras, IoT devices, and other




Figure 3. Query processing flow with network connectivity

high-volume sources, the Edge should process data locally and send only events, alerts, metadata, or necessary segments to the center.

Edge data must also have a clear status. Users need to know whether a displayed result has just been verified online, is valid cached data, or is offline data that requires further verification. Without a visible status, users may mistake a temporary copy for official data.

Cache permission and TTL are determined by LDOP policy as a function of data sensitivity, volatility, freshness requirement, access frequency, legal criticality, impact of stale use, and the ability to reverify the source. High-sensitivity, rapidly changing, or legally critical data should use only minimum attributes or a short-lived reference token and require online verification. Lower-risk, stable, frequently read data may be cached for a bounded TTL.

Operationally, valid cached data may be used only for approved purposes. Expired data must be blocked for legally or operationally important decisions; a warning alone is sufficient only for non-critical informational use explicitly allowed by policy. Pending transactions remain provisional and must not be used as authoritative input to dependent operations. Conflicts may be resolved only by an authorized operational role defined by LDOP policy.

B. Processing flow with network connectivity

During normal operation, a user submits a request through an operational application. The system checks the user's identity, authorization, and purpose. If suitable data is available at the Edge, remains valid, and may be accessed by the user, the Edge can return it immediately. The response must include the data source, update time, and validity period.

If the data is unavailable at the Edge, has expired, or requires online verification, the request is forwarded to the source system through LDOP and the appropriate connection components. After receiving the response, LDOP decides whether part of the result may be cached at the Edge. This decision depends on data policy, sensitivity, frequency of use, and operational needs.

As shown in Figure 3, the flow has two controlled outcomes. A valid and authorized cache hit is returned immediately with its source, version, update time, and validity status. This flow balances speed and control. Simple and frequently repeated requests can be answered quickly at the Edge. Important requests or those requiring the latest data are still verified against the source system.

C. Processing flow during a wan outage

A WAN outage does not have to stop the entire system if valid data remains available at the Edge. The system enters a controlled degraded mode: it performs only approved operations and clearly informs users of the applicable limits.

In controlled degraded mode, the Edge may allow access to data whose time-to-live (TTL) has not expired, display local analytical results, receive new data from devices or operational applications, and place selected transactions in a




Figure 4. Operating principles during a WAN outage

queue. Each transaction should include an event ID, timestamp, user identity, signature or authentication mark, checksum, and a “pending synchronization” status.

Figure 4 shows how the system narrows its functions after loss of WAN connectivity. When connectivity is restored, the Edge sends queued data to LDOP. The central system checks for duplicates, version differences, conflicts, and transaction validity. Valid data is recorded in committed only to the authoritative destination designated by policy; any other repository stores only a non-authoritative copy. If a conflict is detected, the transaction is sent for operational review; it must not automatically overwrite the original data.




Figure 5. Transaction state machine for synchronization and conflict handling in the LDOP-Edge architecture

Figure 5 shows the transaction state machine. Each transaction carries six fields: a source ID, a version, a timestamp or sequence number, an event ID (idempotency key), a checksum, and a signature or MAC. LDOP rejects a transaction in three cases: the authentication mark is invalid, the event ID is a duplicate, or the version is stale. Once LDOP accepts a transaction, it processes the transaction in the order set by the source system. The Edge never resolves a conflict on its own. It never overwrites authoritative data.

Reference tokens follow four rules: they are short-lived, scoped to one purpose, revocable, and they require an online check before any critical operation.

D. Data governance and security

The LDOP-Edge architecture can be deployed safely only when supported by data governance. Before data is sent to an Edge Node, the responsible agency must clearly identify the data owner, data manager, Edge operator, authorized users, purposes of use, retention period, and deletion method.

From a technical perspective, an Edge Node should be protected as part of critical data infrastructure. The system should provide mutual authentication between the Edge and LDOP, encryption in transit and at rest, centralized key management, software integrity checks, patch management, database query monitoring, and centralized collection of security logs.

Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) may be combined. RBAC assigns permissions by role, while ABAC uses contextual attributes. Together, they provide more flexible control than account-based permissions alone.

The threat model assumes that an attacker may compromise or impersonate an Edge/LDOP service; alter, replay, or roll back cached records;

inject or replay offline transactions; steal keys; delete or modify logs; abuse degraded mode or stale data; or attack synchronization after WAN recovery. The corresponding controls are mutual TLS and node certificates; signed or MAC-protected versioned records; monotonic sequence checks and idempotency keys; centralized key rotation and revocation; append-only, remotely collected logs; TTL and purpose enforcement; and validation plus operational review before authoritative commit.

At deployment level, Edge secrets should be held in a TPM, HSM, or protected secret store; cached data should be encrypted at rest; secure boot and signed software should protect software integrity; patches and configuration should be centrally managed; administrative access should use least-privilege RBAC/ABAC; a compromised node should be isolated and its certificate revoked; and expired cache should be removed by cryptographic erasure or an approved secure-deletion process. LDOP is responsible for policy, trust, and centralized evidence, while the Edge enforcement component applies these controls locally.

IV. PROPOSED VALIDATION METHOD

This paper is an architecture and design study and does not report measured performance results. A controlled proof-of-concept is therefore proposed as the next validation stage. It should compare three configurations: a centralized model without an Edge; an LDOP model with a simple cache; and the proposed LDOP-Edge model with policy-controlled caching, an offline queue, and event-based synchronization. All configurations must use the same dataset, workload, network conditions, and operational scenarios.

For reproducibility, the experiment should report Edge and server hardware, software versions, database and middleware, network topology, baseline RTT and bandwidth, outage method and duration, dataset size, read/write ratio, concurrent users, request/response size, TTL policy, number of repetitions, and the calculation of means and variability. The scenarios should include frequent queries,

expired data, local updates, high-volume device data, a defined WAN outage, and synchronization after recovery. Metrics should include mean and P95 latency, queries per second, WAN traffic, cache-hit rate, synchronization and recovery time, duplicates, conflicts, and error rate, together with access-control and log-completeness checks.

TABLE VI. RECOMMENDED METRICS FOR ARCHITECTURE TESTING

Metric group	Specific metrics	Purpose
Performance	Mean latency, P95 latency, and queries per second	Measures responsiveness
Bandwidth	WAN traffic before and after Edge deployment	Measures network-load reduction
Cache	Cache-hit rate and number of TTL expirations	Evaluates caching policy
Synchronization	Synchronization time, number of conflicts, and duplicate transactions	Evaluates reliability after connectivity returns
Security	Denied requests, abnormal access events, and missing logs	Evaluates access control and monitoring
Operations	Recovery time, Edge storage use, and error rate	Evaluates operational continuity

V. DISCUSSION, LIMITATIONS, AND FUTURE WORK

The proposed architecture adds a controlled local-processing layer to existing e-government systems. It is most useful when applications require rapid responses, generate large volumes of data near users, or can safely use short-lived cached data. It is not suitable for every operation. Data with high legal significance, frequent changes, or a strict requirement for the latest status must still be verified against the authoritative source.

The main limitation is increased operational complexity. An agency must manage additional

Edge infrastructure, policies, cryptographic keys, logs, security controls, synchronization processes, and monitoring tools. Poor governance may create inconsistent copies or expand the attack surface. For this reason, Edge deployment should begin with a limited pilot, a clearly defined data scope, and measurable acceptance criteria.

Future work should develop a formal method for selecting data that may be placed at the Edge. It should also study dynamic TTL values based on data volatility, data-risk scoring, conflict-resolution rules, and cost models. A prototype and controlled experiments are needed to quantify improvements in latency, bandwidth use, availability, security, and operating cost.

VI. CONCLUSION

This paper proposes an LDOP–Edge architecture for distributed data processing and sharing in e-government. Its central principle is simple: LDOP manages policies and coordination, source systems retain authority over official data, and Edge Nodes provide only controlled caching and local processing. Each Edge record has a source, version, validity period, and verification status. Offline transactions remain provisional until they are synchronized and reconciled. All access is authenticated, authorized, and logged. The architecture is designed to reduce latency and WAN traffic and may maintain limited services during outages. It may also support local processing of IoT, camera, and AI data. These expected benefits must be measured through a controlled pilot and must not weaken data governance, protection, or system security.

REFERENCES

- [1] Government of Vietnam, Decision No. 2439/QĐ-TTg, National Data Architecture Framework, National Data Governance and Management Framework, and Shared Data Dictionary (Version 1.0), Nov. 4, 2025.
- [2] National Assembly of Vietnam, Data Law, Nov. 30, 2024.
- [3] Government of Vietnam, Decree No. 194/2025/ND-CP, July 3, 2025.
- [4] Government of Vietnam, Decree No. 278/2025/ND-CP, Oct. 22, 2025.
- [5] Government of Vietnam, Decree No. 47/2020/ND-CP, Apr. 9, 2020.
- [6] M. Satyanarayanan, “The Emergence of Edge Computing,” *Computer*, vol. 50, no. 1, pp. 30-39, 2017.
- [7] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge Computing: Vision and Challenges,” *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637-646, 2016.
- [8] M. Iorga et al., Fog Computing Conceptual Model, NIST SP 500-325, 2018.
- [9] V. C. Hu et al., Guide to Attribute Based Access Control (ABAC) Definition and Considerations, NIST SP 800-162, 2019.
- [10] S. Rose et al., Zero Trust Architecture, NIST SP 800-207, 2020.
- [11] OECD, The Path to Becoming a Data-Driven Public Sector, OECD Digital Government Studies, 2019.
- [12] J. Haas, R. Mogk, E. Yanakieva, A. Bieniusa, and M. Mezini, “LoRe: A Programming Model for Verifiably Safe Local-First Software”, *ACM Transactions on Programming Languages and Systems*, vol. 46, no. 1, Article 2, 2024, DOI: 10.1145/3633769.
- [13] K. Jannes, B. Lagaisse, and W. Joosen, “Secure Replication for Client-centric Data Stores”, *Proc. 3rd Int. Workshop on Distributed Infrastructure for the Common Good*, pp. 31-36, 2022, DOI: 10.1145/3565383.3566111.
- [14] A. Jeffery and R. Mortier, “AMC: Towards Trustworthy and Explorable CRDT Applications with the Automerge Model Checker”, *Proc. 10th Workshop on Principles and Practice of Consistency for Distributed Data*, 2023, DOI: 10.1145/3578358.3591326.
- [15] E. Bozkaya-Aras et al., “Blockchain-Based Secure Content Caching and Computation in Edge Computing Systems”, *IEEE Access*, 2024, DOI: 10.1109/ACCESS.2024.3383505.
- [16] H. T. Ho, N. T. Hung, and N. T. Minh, “Using Distributed Computer Networks to Build Block Cryptanalysis Solutions”, *Journal of Science and Technology on Information Security*, vol. 1, no. 13, pp. 71-79, 2022, DOI: 10.54654/isj.v1i13.232.
- [17] N. Q. Khanh, T. Q. Duc, N. V. Toan, and T. V. Van, “FedPC Cloud CPU Forecasting Using Federated Learning and Periodicity-Based Clustering”, *Journal of Science and Technology on Information Security*, vol. 2, no. 25, pp. 52-60, 2025, DOI: 10.54654/isj.v2i25.1115.

ABOUT THE AUTHOR

Pham Van Dinh



Workplace: Department of Information Technology, Ministry of Public Security, Vietnam

Email: phamdinhbca@gmail.com

Education and professional experience: He graduated from the People's Security Academy in 2001.

He received his M.S. degree in Information Technology from the VNU University of Engineering and Technology in 2009 and his Ph.D. degree in Social Order and Security from the People's Police Academy in 2019.

Research interests: His research and professional interests include information technology, digital transformation, and professional technical solutions for public security operations.

Tên tác giả: **Phạm Văn Dinh**

Cơ quan công tác: Cục Công nghệ thông tin, Bộ Công an, Việt Nam

Email: phamdinhbca@gmail.com

Quá trình đào tạo và công tác: Tốt nghiệp Học viện An ninh nhân dân năm 2001; nhận bằng Thạc sĩ Công nghệ thông tin tại Trường Đại học Công nghệ, Đại học Quốc gia Hà Nội năm 2009; nhận bằng Tiến sĩ chuyên ngành An ninh và trật tự xã hội tại Học viện Cảnh sát nhân dân năm 2019.

Hướng nghiên cứu hiện nay: Nghiên cứu, ứng dụng công nghệ thông tin, chuyển đổi số và các giải pháp kỹ thuật phục vụ công tác nghiệp vụ Công an

Ha Mai Huyen Trang



Workplace: Department of Information Technology, Ministry of Public Security, Vietnam

Email: 23078027@vnu.edu.vn

Education: She graduated from the People's Security Academy in 2014 and received her M.S. degree in

Information Technology from the Posts and Telecommunications Institute of Technology in 2018. She is currently a Ph.D. candidate at VNU International School, Vietnam National University, Hanoi.

Research interests: Her research interests include machine learning, federated learning, digital transformation, and e-government.

Tên tác giả: **Hà Mai Huyền Trang**

Cơ quan công tác: Cục Công nghệ thông tin, Bộ Công an, Việt Nam

Email: 23078027@vnu.edu.vn

Quá trình đào tạo: Tốt nghiệp Học viện An ninh nhân dân năm 2014; nhận bằng Thạc sĩ Công nghệ thông tin tại Học viện Công nghệ Bưu chính Viễn thông năm

2018. Hiện là nghiên cứu sinh tại Trường Quốc tế, Đại học Quốc gia Hà Nội.

Hướng nghiên cứu hiện nay: Học máy, học liên kết, chuyển đổi số và chính phủ điện tử.



Pham Thi Son

Workplace: Department of Information Technology, Ministry of Public Security, Vietnam

Email: phamsonbca@gmail.com

Education: She received her Engineer's degree in Information

Technology from Hung Yen University of Technology and Education in 2007.

Research interests: Her research and professional interests include information technology applications, digital transformation, and technical solutions supporting public security operations.

Tên tác giả: **Phạm Thị Sơn**

Cơ quan công tác: Cục Công nghệ thông tin, Bộ Công an, Việt Nam

Email: phamsonbca@gmail.com

Quá trình đào tạo: Tốt nghiệp kỹ sư ngành Công nghệ thông tin tại Trường Đại học Sư phạm Kỹ thuật Hưng Yên năm 2007.

Hướng nghiên cứu hiện nay: Nghiên cứu, ứng dụng công nghệ thông tin, chuyển đổi số và các giải pháp kỹ thuật phục vụ công tác nghiệp vụ Công an.

Vu Ngoc Duong



Workplace: GTSC Vietnam technology services and commercial joint stock company

Email: duongvn.bca@gmail.com

Education: He received the B.Sc. degree in 2021 from Hanoi

University of Industry, and the M.Sc. degree in 2026 from Posts and Telecommunications Institute of Technology.

Recent research direction: His research interests include intelligent information systems, network optimization, large-scale data communication, and AI applications in networking.

Tên tác giả: **Vũ Ngọc Dương**

Cơ quan công tác: Công ty cổ phần thương mại và dịch vụ công nghệ GTSC Việt Nam

Email: duongvn.bca@gmail.com

Quá trình đào tạo: Tốt nghiệp Cử nhân năm 2021 tại trường Đại học Công nghệ Hà Nội, Thạc sĩ năm 2026 tại trường Học viện Công nghệ Bưu chính Viễn Thông.

Hướng nghiên cứu hiện nay: Hệ thống thông tin thông minh, tối ưu hóa mạng, truyền thông dữ liệu quy mô lớn và ứng dụng trí tuệ nhân tạo trong mạng máy tính.