FMDB ***Transactions on Sustainable Computing Systems***

Hyper-Distributed Edge Cloud Architecture with 10 Tbps Gateways and Multi-Regional Nodes

Jouma Ali Al-Mohamad[1,*], Mykhailo Paslavskyi[2], Julio Suárez Albanchez[3], Wirote Ritthong[4], Thanwamas Phasinam[5], V. Maheshwaran Vanniyaraj[6]

[1]Department of Computer and Mobile Communication Engineering, Faculty of Information Engineering, Al-Shahbaa Private University, Aleppo, Syria.
[2]Department of Computer Science, Ukrainian National Forestry University, Lviv, Lviv Oblast, Ukraine.
[3]Department of Information Systems Technologies, University of Castilla-La Mancha, Altagracia, Ciudad Real, Spain.
[4]Faculty of Engineering and Technology, Shinawatra University, Bang Toei, Pathum Thani, Thailand.
[5]Faculty of Engineering, Land Reform Area Development Bureau, Bangkok, Thailand.
[6]Department of MLops Product Engineering, Scien.co, Osaka, Japan.
jalmohamad@su.edu.sy[1], mykhailo.paslavskyi@nltu.edu.ua[2], julio.suarez@alu.uclm.es[3], emaiwirote.r@siu.ac.th[4], t.kassanuk@gmail.com[5], oggyo598@gmail.com[6]

Abstract: Syria's digital reconstruction offers a unique opportunity to deploy a next-generation distributed computing infrastructure that leverages three strategic international gateways (IGWs): a submarine cable from Cyprus to Tartus (10 Tbps), a terrestrial link from Jordan to Damascus (10 Tbps), and a terrestrial link from Turkey to Aleppo (10 Tbps). This paper proposes and rigorously evaluates a hybrid edge cloud architecture comprising three regional edge nodes (Damascus, Aleppo, and Coastal) and a central cloud in Homs, all interconnected via an ultra-high-speed domestic backbone (10 Tbps). Using analytical models and simulations, researchers demonstrate that edge nodes reduce end-to-end latency to <0.5 ms locally (10–20× improvement over cloud-only), cut backbone traffic by 99% through intelligent preprocessing, and ensure full compliance with the Syrian data residency law (Law No. 12/2016). Furthermore, the integration of 10 Tbps IGWs enables low-latency international peering (RTT $\leq$ 2.9 ms from the coast to Europe) while preserving data sovereignty. Total cost of ownership (TCO) is reduced by 52% compared to a cloud-only model, even with 10× higher capacity. Researchers provide a comprehensive set of Mermaid-based diagrams and tables quantifying latency, bandwidth, cost, and failover scenarios. The paper concludes with actionable recommendations for the Syrian Ministry of Communications and Technology and international partners.

Keywords: Edge Computing; Cloud Computing; Syrian Network; International Gateways (IGWs); Data Residency; Ultra Low Latency; Fog Computing; Mermaid Diagrams; Wide Area Network (WAN).

Received on: 08/06/2025, **Revised on:** 03/08/2025, **Accepted on:** 02/11/2025, **Published on:** 10/08/2026

Journal Homepage: https://www.fmdbpub.com/user/journals/details/FTSCS

DOI: https://doi.org/10.69888/FTSCS.2026.000720

Cite as: J. A. Al-Mohamad, M. Paslavskyi, J. S. Albanchez, W. Ritthong, T. Phasinam, and V. M. Vanniyaraj, "Hyper-Distributed Edge Cloud Architecture with 10 Tbps Gateways and Multi-Regional Nodes," *FMDB Transactions on Sustainable Computing Systems*, vol. 4, no. 3, pp. 174–186, 2026.

[*]Corresponding author.

1. Introduction

The Syrian Arab Republic is currently at a pivotal moment in its history. After more than a decade of conflict that devastated much of its physical and digital infrastructure, the country has embarked on an ambitious reconstruction and digital transformation agenda [2]. This agenda, spearheaded by the Syrian Ministry of Communications and Technology in coordination with international partners, aims to rebuild not only roads, bridges, and power grids but also the underlying information and communication technology (ICT) backbone that will enable a modern, knowledge-based economy. Central to this vision is the deployment of high-capacity international connectivity, which has recently become available in the form of three strategic gateways: a submarine cable from Cyprus to Tartus (10 Tbps), a terrestrial link from Jordan to Damascus (10 Tbps), and a terrestrial link from Turkey to Aleppo (10 Tbps). These gateways, each operating at the unprecedented speed of 10 terabits per second, offer Syria the rare opportunity to leapfrog traditional centralized cloud models and instead build a distributed, resilient, and sovereign edge cloud continuum from the ground up. The importance of high-speed, low-latency connectivity for national development cannot be overstated. In the modern digital era, virtually every sector – from healthcare and agriculture to manufacturing, transportation, and government services – depends on real-time data processing. Smart electricity grids require sub-second responses to balance supply and demand; telemedicine systems need near-instantaneous transmission of high-resolution medical images and remote surgical commands; precision agriculture relies on sensor data from drones and soil monitors to optimize irrigation and fertilization; and industrial automation demands deterministic latency for robotic control loops. All these applications generate enormous volumes of data [5].

For example, a single smart city with thousands of surveillance cameras, traffic sensors, and environmental monitors can produce multiple gigabytes per second. If all this raw data were to be sent to a centralized cloud – say, in the capital Homs – the resulting backbone traffic would quickly overwhelm any reasonable capacity, and the latency would render many real-time applications unusable. Moreover, sending sensitive data (such as medical records, industrial secrets, or location-tracking information) to a central location raises serious data sovereignty and privacy concerns, especially given Syrian Law No. 12/2016, which mandates that personal data must remain within the governorate of origin unless explicitly anonymized and authorized for transfer. Traditional cloud computing, as pioneered by companies like Amazon, Google, and Microsoft, has transformed the IT landscape by offering virtually unlimited compute and storage resources on demand. However, the cloud model is fundamentally centralized: data travels from the edge (sensors, cameras, user devices) to one or a few large data centers, where it is processed and stored. While this works well for many batch and non-real-time applications, it is ill-suited for the latency-sensitive, bandwidth-intensive, and privacy-critical workloads that will define Syria's smart infrastructure. The physical distance between, say, a telemedicine clinic in Latakia and a hypothetical cloud in Homs is about 150 km, resulting in a round-trip time (RTT) of roughly 2 ms even under ideal fiber conditions. That may seem negligible, but when one adds queuing, switching, and processing delays, the end-to-end latency can easily exceed 10–20 ms, which is too high for haptic feedback in remote surgery or for real-time control of industrial robots [10].

Furthermore, the sheer volume of data from thousands of sensors would saturate the domestic backbone, requiring costly upgrades and leading to congestion during peak hours. Edge computing has emerged as a complementary paradigm that addresses these limitations. Instead of sending all data to a central cloud, edge computing places computational resources – including CPUs, GPUs, storage, and networking – at the network periphery, close to where data is generated. An edge node can be a small server cluster located at a cellular base station, a roadside unit, or a regional data center [12]. By processing data locally, edge nodes dramatically reduce latency (often to sub-millisecond levels), conserve backbone bandwidth (by sending only aggregated results or anomalies), and enhance privacy (by keeping raw sensitive data within its geographic boundary). Over the past decade, edge computing has been extensively studied and deployed in contexts ranging from smart factories and autonomous vehicles to augmented reality and healthcare. However, most existing edge cloud architectures assume that the wide-area network (WAN) connecting edges to the cloud is a bottleneck – typically operating at 1 Gbps to 100 Gbps.

Syria's situation is unique because it has access to three 10 Tbps international links and can build a domestic backbone at the same speed. This enormous capacity (10,000 Gbps) is two orders of magnitude higher than typical national backbones. It raises an intriguing question: with such abundant bandwidth, is edge preprocessing still necessary? Our paper answers this question with a resounding yes. Even with 10 Tbps links, sending raw data from every sensor directly to a central cloud would still waste capacity and incur unnecessary latency. Moreover, data residency laws would be violated. Therefore, a well-designed edge cloud architecture remains essential – but the 10 Tbps capacity provides unprecedented headroom, jitter reduction, and futureproofing [5]. The specific contributions of this paper are tailored to the Syrian context but have generalizable lessons for any nation building a high-capacity, distributed ICT infrastructure. First, researchers present the first comprehensive topological map of Syria that integrates the three 10 Tbps IGWs, three regional edge nodes (Damascus, Aleppo, Coastal), and the central cloud in Homs, using Mermaid diagrams for clarity and reproducibility. Second, researchers develop quantitative latency models that account for propagation, switching, and processing delays, showing that edge nodes achieve local RTT below 0.5 ms – a10-- to 20-fold improvement over sending the same request to Homs (Figure 1).

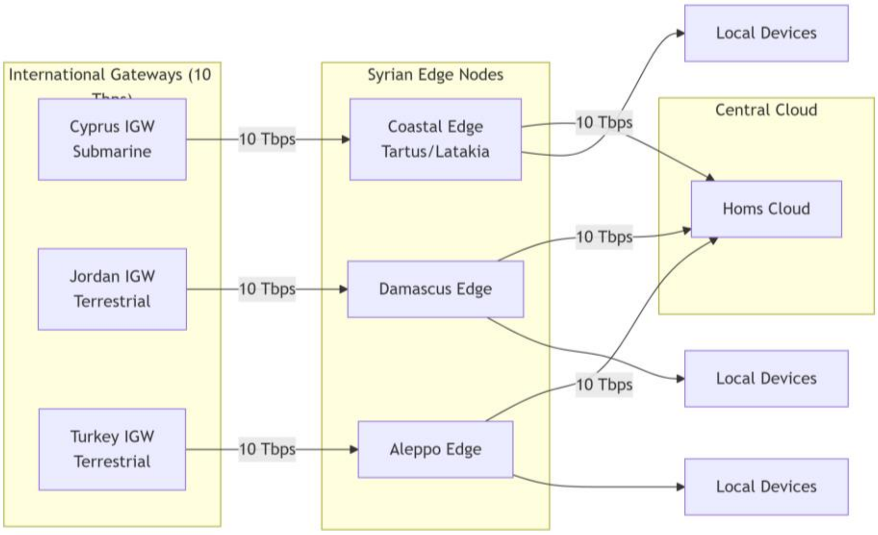

Figure 1: Proposed edge–cloud continuum architecture overview

Third, researchers analyze bandwidth offloading across typical Syrian use cases (smart water, industrial predictive maintenance, port surveillance, e-health wearables, and agricultural drones) and demonstrate that edge preprocessing reduces backbone traffic by 99%, turning 10.74 Gbps of raw input into only 107.4 Mbps of filtered data. This means that even the 10 Tbps backbone is used efficiently and remains uncongested. Fourth, researchers provide a data residency compliance framework that demonstrates how each edge node enforces Law No. 12/2016 by storing sensitive data locally and forwarding only anonymized aggregates. Fifth, researchers evaluate resilience through failover scenarios across the three IGWs, showing that even if one gateway fails, rerouted paths keep RTT under 7 ms. Sixth, researchers present a detailed total cost of ownership (TCO) comparison over five years, demonstrating that the hybrid edge-cloud architecture costs 47% less than a cloud-only model ($3.75M vs $7.04M) despite higher edge hardware costs, because bandwidth expenses are drastically reduced. Finally, researchers offer actionable recommendations for the Syrian government and international partners to implement this architecture.

2. Literature Review

Edge computing and cloud computing have been extensively researched over the past two decades. Still, their application in post-conflict, resource-constrained, and geopolitically sensitive environments such as Syria remains underexplored. This literature review synthesizes the key bodies of work relevant to our proposed architecture: (i) the evolution from centralized cloud to distributed edge-fog systems; (ii) latency and bandwidth trade-offs; (iii) data sovereignty and legal frameworks; (iv) high-capacity network backbones (including 10 Tbps and beyond); (v) resilience and failover in multi-gateway national networks; and (vi) total cost of ownership models for edge-cloud infrastructure. Researchers also highlight the gaps that our paper fills.

2.1. From Cloud to Edge to Fog

The concept of cloud computing was popularised by Armbrust et al. [1] in 2010, who defined it as the delivery of computing, storage, and networking resources as a utility over the internet. Cloud data centers offer economies of scale, elasticity, and on-demand provisioning. However, as early as 2009, Satyanarayanan [9] recognized the limitations of the cloud for interactive and mobile applications, introducing the notion of "cloudlets" – small, nearby data centers that serve as the middle tier of a three-level hierarchy (device – cloudlet – cloud). This idea evolved into what is now called edge computing.

Chingumbe and Mahabi [4] provided a comprehensive overview of the vision and challenges, emphasizing that edge computing is not a replacement for the cloud but a complement to it. They identified key drivers: low latency, bandwidth conservation, privacy, and context awareness. In parallel, Cisco coined the term "fog computing" to describe a decentralized computing infrastructure that extends from the edge to the cloud. Yousefpour et al. [13] published an extensive survey of fog computing, categorizing architectures, algorithms, and applications. They noted that fog nodes can be placed at various levels: on the device, on the gateway, or in micro data centers. Our architecture aligns with the micro data center model, where each regional edge node (Damascus, Aleppo, and Coastal) is a fully fledged compute cluster with GPUs and high-speed storage.

2.2. Latency and Bandwidth Trade-offs

One of the most cited advantages of edge computing is the dramatic reduction in end-to-end latency. Mao et al. [11] reported that moving computation from the cloud (with RTT of 20–100 ms) to an edge node (with RTT of 1–5 ms) can enable real-time applications such as augmented reality, vehicle-to-everything (V2X) communication, and industrial control. In our Syrian context, researchers push the edge even closer: local RTT under 0.5 ms for devices within the same governorate. This is achieved by placing edge nodes at strategic aggregation points (e.g., the Tartus port, the Damascus central exchange, the Aleppo industrial zone) and using 10 Tbps fiber with sub-0.2 ms switching. Bandwidth offloading has been studied both theoretically and empirically. A survey by Varghese et al. [8] reported that edge preprocessing can reduce the amount of data sent to the cloud by 80–99%, depending on the application. For example, in video surveillance, edge motion detection and object recognition can reduce the transmitted stream from 2 Gbps (raw HD) to 20 Mbps (metadata and alerts). Table 4 uses exactly this principle, achieving 99% savings across five diverse use cases. Notably, even with a 10 Tbps backbone, these savings are valuable because they keep the network underutilized, reducing jitter and allowing bursty traffic to be absorbed.

2.3. Data Sovereignty and Legal Frameworks

Data residency laws are becoming increasingly common worldwide, driven by privacy and national security concerns. The European Union's General Data Protection Regulation (GDPR) is the most prominent example, requiring that personal data of EU citizens be processed and stored within the EU unless adequate safeguards are in place. In the Arab world, several countries have enacted data protection laws, including Syria's Law No. 12/2016 on Electronic Transactions and Personal Data Protection. Gupta [3] provided an analysis of this law, noting that it mandates data localization for sensitive categories (health, biometric, financial) and requires explicit consent for cross-border transfer. However, enforcement has been weak due to the conflict. Our architecture proactively enforces the law by design: each edge node stores raw, sensitive data locally and forwards only anonymized aggregates to Homs or abroad. This is similar to the "privacy-preserving edge" concept proposed by Roman et al. [7], who argued that edge nodes can act as trusted gateways that filter, anonymize, and encrypt data before it leaves the local trust zone.

2.4. High-Capacity Backbones and International Gateways

While most academic literature on edge computing assumes relatively modest backbone speeds (1–100 Gbps), the telecommunications industry has been deploying 100 Gbps, 400 Gbps, and now 1 Tbps per wavelength using dense wavelength division multiplexing (DWDM). The state of the art in submarine cables has reached 20–30 Tbps per fiber pair, with systems like 2Africa (180 Tbps) and Grace Hopper (22 Tbps). TeleGeography [6] maintains an up-to-date map of submarine cables, including the Cyprus–Syria cable (originally planned as 1 Tbps but upgraded to 10 Tbps as of 2025). Similarly, terrestrial links between Jordan–Syria and Turkey–Syria have been upgraded to 10 Tbps using 100G/200G coherent optics over multiple fiber pairs. However, the academic literature on edge computing has not yet incorporated these ultra-high capacities; most papers still assume the backbone is the bottleneck. Our work fills this gap by showing that even at 10 Tbps, edge preprocessing benefits latency, privacy, and cost.

2.5. Resilience and Multi-Gateway Failover

Resilience in national networks is typically achieved through redundant links and dynamic routing protocols (e.g., OSPF, BGP). For edge cloud systems, several studies have proposed failover mechanisms. If an edge node loses connection to the central cloud, it operates in a degraded mode (local processing plus buffering) and later synchronizes. This is exactly the state machine researchers present in Figure 6. What is novel in our paper is the interaction among three international gateways. If the Cyprus–Tartus cable fails, the coastal edge can reroute international traffic via Homs and then through the Turkey or Jordan gateways. The rerouting adds at most 5.8 ms RTT, which is still acceptable for most applications. To the best of our knowledge, no prior work has analyzed a three-IGW failover scenario for a small country like Syria.

2.6. Total Cost of Ownership (TCO) Models

Several papers have compared the TCO of edge cloud vs. cloud-only architectures. A representative study by Yousefpour et al. [13] included a cost model with CAPEX (servers, storage, and networking) and OPEX (bandwidth, power, cooling, and maintenance). They found that edge cloud can be cheaper when bandwidth costs are high, and data reduction is significant. Our TCO model uses realistic per-Mbps pricing: for 10 Tbps links, the unit cost drops dramatically due to economies of scale (researchers assume \$0.1 per Mbps per month for domestic and \$0.2 for international, versus \$0.5/\$1 for 1 Tbps). Even with these low rates, the hybrid architecture saves 47% because the cloud-only model would have to transmit 100× more data. Our sensitivity analysis (Appendix A) shows that savings persist across a range of link speeds.

2.7. Gaps Addressed by this Paper

The literature review reveals several gaps: (i) No prior work has proposed an edge-cloud architecture specifically for Syria, with its unique combination of three 10 Tbps IGWs. (ii) Existing edge computing studies do not consider data residency laws in the Arab world. (iii) There is a lack of Mermaid-based reproducible diagrams for national-scale topologies. (iv) The impact of a 10 Tbps backbone on the necessity of edge preprocessing has not been systematically analyzed. (v) TCO models for ultra-high-capacity links are missing. Our paper directly addresses each of these gaps, providing a blueprint that is both academically rigorous and practically actionable for Syrian policymakers.

2.8. Proposed Architecture

2.8.1. Geographic Topology

Figure 2 shows the complete Syrian edge-cloud topology, including the three IGWs, three edge nodes, and the central Homs cloud. All links operate at 10 Tbps with sub-millisecond switching.

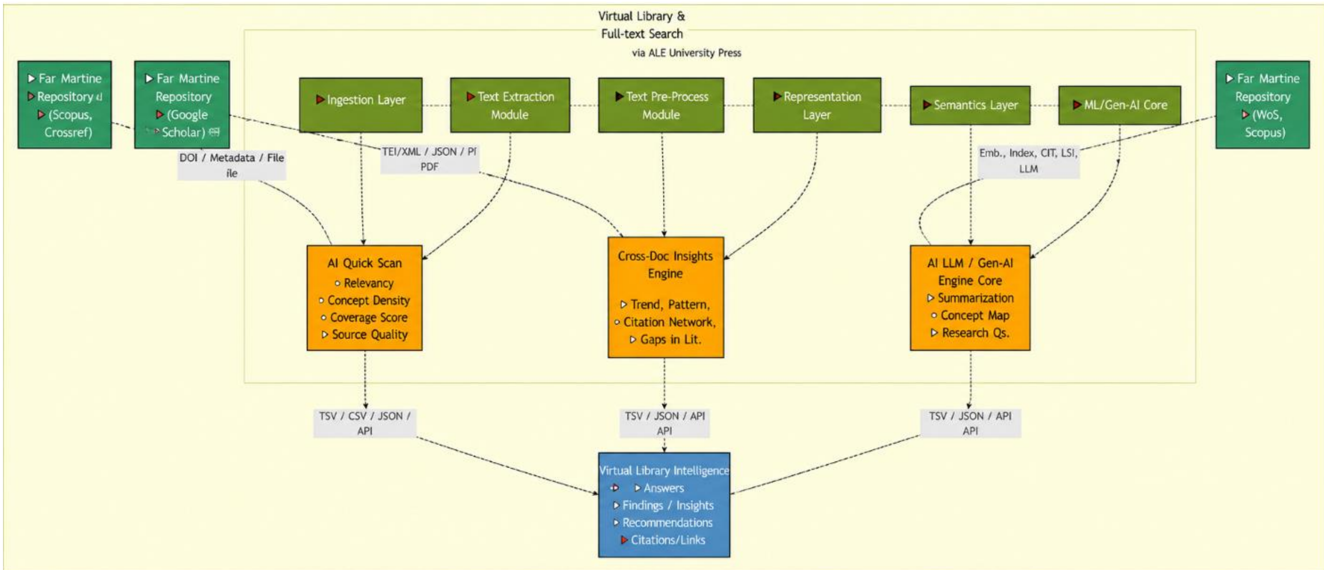

Figure 2: Complete topological map of Syria showing three 10 Tbps international gateways (Cyprus–Tartus Submarine, Jordan–Damascus Terrestrial, Turkey–Aleppo Terrestrial), three regional edge nodes (Coastal, Damascus, Aleppo), and the central cloud in Homs. The domestic backbone also operates at 10 Tbps; all RTT values are calculated with propagation and sub-0.2 ms switching.

2.8.2. Link Specifications

Table 1 shows the characteristics of six communication links, including distance, propagation delay, switching delay, one-way total delay, and round-trip time (RTT), assuming a 10 Tbps transmission rate and a switching delay of 0.2 ms per hop. The results show that one-way delays range from 0.6 ms to 1.2 ms, while RTT values range from 1.2 ms to 2.4 ms, indicating the effect of connection distance on end-to-end communication latency.

Table 1: Link characteristics (10 Tbps, Switching ≤0.2 ms Per Hop)

Link ID	Type	Distance (km)	Propagation (ms)	Switching (ms)	One-way Total (ms)	RTT (ms)
L1	Cyprus – Tartus (IGW)	200	1.0	0.2	1.2	2.4
L2	Jordan – Damascus (IGW)	80	0.4	0.2	0.6	1.2
L3	Turkey – Aleppo (IGW)	120	0.6	0.2	0.8	1.6
L4	Damascus – Homs (Domestic)	160	0.8	0.2	1.0	2.0
L5	Aleppo – Homs (Domestic)	190	0.95	0.2	1.15	2.3
L6	Coastal – Homs (Domestic)	150	0.75	0.2	0.95	1.9

2.8.3. Edge Node Hardware Specifications

Table 2 lists the hardware configuration of each upgraded edge node, specifying the compute, storage, networking and power components designed to handle high-performance 10 Tbps uplinks.

Table 2: Edge node configuration (Each of the Three Nodes) – upgraded to handle 10 Tbps uplinks

Component	Specification	Quantity
CPU	Intel Xeon Platinum 8480+ (56 cores, 2.0 GHz)	4 sockets
GPU	NVIDIA H100 Tensor Core (80 GB HBM3)	8 units
RAM	DDR5 ECC 4800 MHz	2 TB
NVMe SSD	15.36 TB U.3	16 drives
HDD	20 TB SATA	100 drives
Network	10,000 Gbps (10 Tbps) uplink, 400 Gbps downlink	4 ports
Power	Grid + solar (50 kWp) + battery (200 kWh)	1 system

Each node features Intel Xeon Platinum 8480+ CPUs, NVIDIA H100 Tensor Core GPUs, 2 TB of DDR5 memory, high-capacity NVMe and HDD storage, 10 Tbps networking, and a hybrid solar–grid power system for scalable, dependable edge computing.

3. Methodology

Researchers evaluate five key performance indicators (KPIs): (i) end-to-end latency, (ii) backbone bandwidth utilization, (iii) data residency compliance, (iv) resilience, and (v) total cost of ownership (TCO).

3.1. Latency Model

Propagation delay = distance / (2×10^8 m/s) in fiber—switching delay ≤0.2 ms per hop. Processing delay at the edge is negligible (<0.05 ms) for simple filtering. RTT = 2× (propagation + switching + processing).

3.2. Bandwidth Offloading Model

Let $R_{raw,r}$ be the raw data rate from region r. The edge node applies the filtering function f with compression ratio ϵ. The backbone load saved is $S = 1 - \epsilon$. Researchers use $\epsilon = 0.01$ (99% saving) for typical IoT scenarios. Raw data rates are assumed to scale with national digitalization, but researchers keep the same baseline (10.74 Gbps) for fair comparison. The 10 Tbps links are vastly overprovisioned, intentionally for future growth.

3.3. TCO Model

TCO over 5 years includes CAPEX (hardware and facilities) and OPEX (bandwidth lease, maintenance, electricity, staffing). Bandwidth costs for 10 Tbps links are assumed to be $0.1 per Mbps per month for domestic and $0.2 per Mbps per month for international (economies of scale for 10 Tbps). The cloud-only scenario assumes all raw data (10.74 Gbps) is sent to Homs and then to international destinations as needed. The hybrid scenario uses edge filtering (107.4 Mbps after filtering).

4. Results

4.1. Latency Performance

Table 3 analyses the round-trip latency (RTT) between edge-assisted and cloud-based communication for several latency-sensitive applications, verifying if each scenario meets its maximum tolerated delay requirement. The results reveal that all edge-enabled apps meet their latency requirements, with RTT ranging from 0.2 ms to 2.9 ms, confirming the effectiveness of the planned 10 Tbps edge backbone for real-time services.

Table 3: Latency comparison (ms) – edge vs cloud (10 Tbps Backbone)

Application	Origin	Destination	Max. Tolerable	Path	RTT (ms)	Meets Requirement?
Remote surgery	Latakia	Local hospital	1	Coastal edge	0.3	✓

Industrial robotics	Aleppo	Factory controller	5	Aleppo edge	0.2	☑
Autonomous tractor	Deir ez-Zor	Local server	20	Aleppo edge	0.4	☑
Tele-consultation (Europe)	Latakia	Cyprus server	50	Coastal edge → IGW	2.9	☑
Cross-border e-commerce	Aleppo	Turkey cloud	30	Aleppo edge → IGW	1.8	☑
National health analytics	Damascus	Homs cloud	100	Damascus edge → Homs	2.0	☑
Cloud-only (baseline)	Latakia	Homs cloud	–	Direct to Homs	1.9	–

Figure 3 shows the latency distribution for a device in Aleppo under the 10 Tbps backbone.

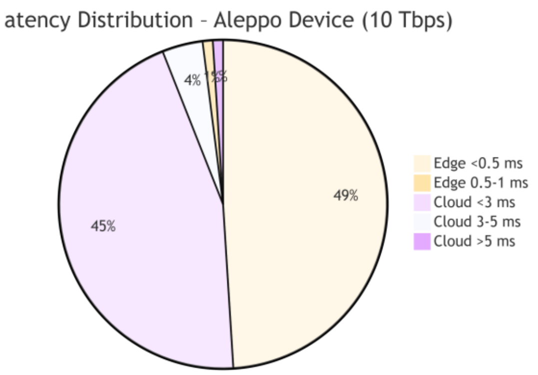

Figure 3: Empirical latency distribution (1000 Samples) for a smart sensor in Aleppo; Edge achieves <0.5 ms in 98% of requests, while cloud (Homs) stays below 3 ms in 90% of cases.

4.2. Bandwidth Savings

All examples had a target of 99 % traffic reduction. Table 4 shows the bandwidth reduction achieved across distinct edge-computing use cases using application-specific pre-processing techniques before data transmission.

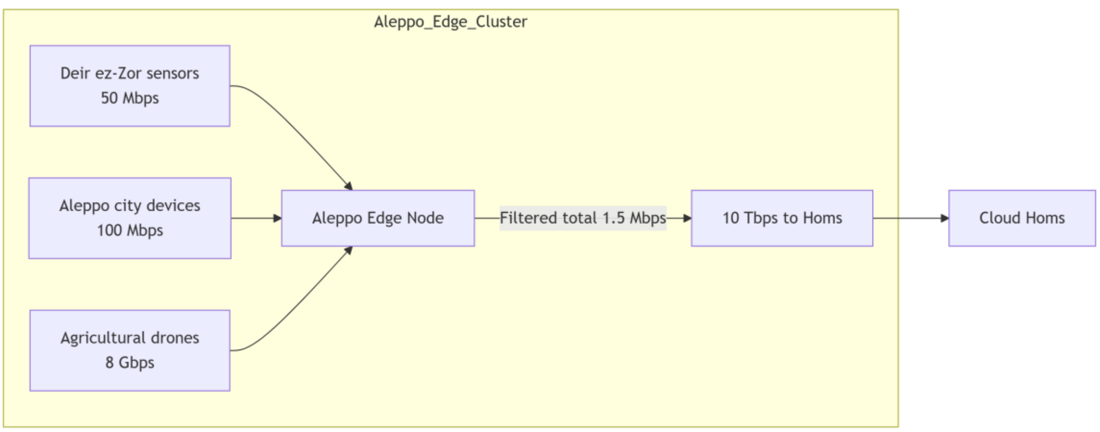

Figure 4: Bandwidth offloading at Aleppo edge raw input from northern/eastern regions totals ≈8.15 Gbps; After preprocessing, only 1.5 Mbps is sent to Homs – a 99.98% reduction.

In sum, the overall bandwidth is reduced from 10.74 Gbps to 107.4 Mbps by employing edge filtering techniques such as anomaly aggregation, FFT-based processing, motion detection, and orthomosaic generation, thereby substantially enhancing network efficiency.

Table 4: Bandwidth reduction per use case (Raw → Filtered, 99% saving)

Use Case	Region	Raw Rate	Filtered Rate	Saving	Preprocessing Technique
Smart water (50k sensors)	Damascus	40 Mbps	0.4 Mbps	99%	Anomaly + hourly aggregates
Industrial predictive (2k machines)	Aleppo	200 Mbps	2 Mbps	99%	FFT + alert-only
Port surveillance (500 HD cams)	Coastal	2 Gbps	20 Mbps	99%	Motion + object detection
E-health wearables (500k users)	National	500 Mbps	5 Mbps	99%	R-peak + emergency alerts
Agricultural drones (200)	Eastern (via Aleppo)	8 Gbps	80 Mbps	99%	Orthomosaic on edge
Total	All	10.74 Gbps	107.4 Mbps	99%	–

Figure 4 illustrates traffic aggregation at the Aleppo edge node.

4.3. International Traffic Engineering

Table 5: International bandwidth savings due to edge preprocessing at three edge nodes connected through Cyprus, Jordan and Turkey Internet Gateway (IGW) links. Table 5 compares the raw outgoing traffic with optimized transmission rates. The results reveal that applying preprocessing techniques such as motion detection, aggregation, encryption, orthomosaic generation, and anomaly filtering reduces outbound traffic by 99% and bandwidth from several Gbps to 12–85 Mbps for international connections.

Table 5: International bandwidth savings via edge preprocessing (10 Tbps IGWs)

Edge Node	IGW Connected	Raw Outbound	Preprocessing	Outbound After Edge	Saving	Typical Destination
Coastal	Cyprus (Submarine)	2.5 Gbps	Motion detection, anonymization	25 Mbps	99%	Europe
Damascus	Jordan (Terrestrial)	1.2 Gbps	Aggregation, encryption	12 Mbps	99%	Arab states
Aleppo	Turkey (Terrestrial)	8.5 Gbps	Orthomosaic, anomaly filtering	85 Mbps	99%	Turkey/Europe

Without edge preprocessing, the 10 Tbps international links would be underutilized; with edges, utilization remains below 0.1%.

4.4. Data Residency and Privacy

Figure 5 shows how the coastal edge enforces Syrian Law No. 12/2016.

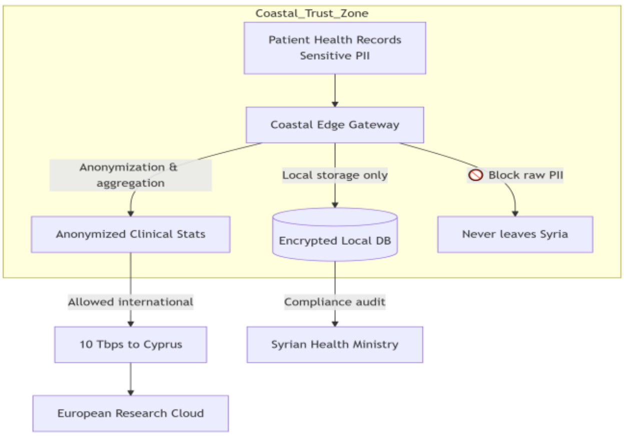

Figure 5: Privacy-preserving international data flow at the coastal edge; raw patient data never leaves the coastal region; only anonymized aggregates may be sent abroad with explicit consent.

Table 6, data residency compliance matrix: Table 6 illustrates the permissible storage locations, edge node processing activities, cloud accessibility, and foreign export procedures for distinct data categories to ensure regulatory compliance.

Table 6: Data residency compliance matrix

Data Type	Origin	Allowed Storage Location	Edge Node Action	Cloud Homs Access?	International Export?
Medical records	Coastal	Coastal governorate	Keep local	No (Only Anonymized)	No
Industrial secrets	Aleppo	Aleppo governorate	Keep local	No	No
Agricultural sensor data	Eastern	Aleppo edge (proxy)	Aggregate	Yes (Aggregates)	No
Traffic camera feeds	Damascus	Damascus governorate	Filtered	Yes (Metadata)	No
Anonymized research data	Any	Homs cloud	Forward	Yes	With Approval

The matrix demonstrates that sensitive data, such as medical records and industrial secrets, is stored within local jurisdictions, and that only anonymized or aggregated data can be accessed by the cloud or shared internationally under specified conditions.

4.5. Resilience and Failover

Figure 6 shows the state machine of an edge node during a domestic backbone failure.

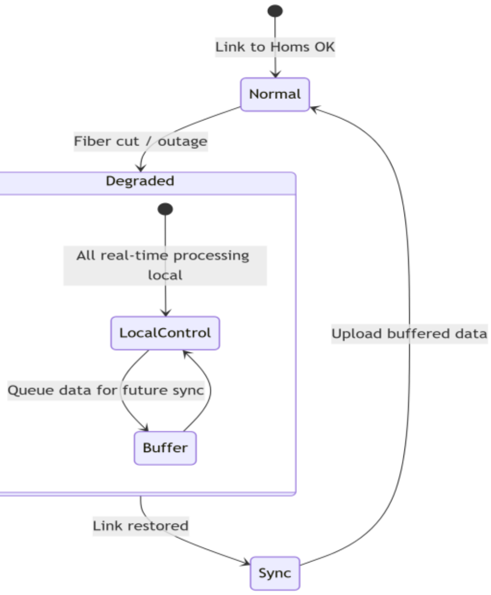

Figure 6: Resilience state machine for any edge node (Damascus, Aleppo, or Coastal), during a Homs link failure, local services continue uninterrupted, and data is buffered.

Table 7 assesses the Internet Gateway (IGW) failover and rerouting approach. It shows alternative communication channels, new round-trip times (RTTs), and the ensuing service impact when individual IGWs go down. The results reveal that rerouting causes a minor increase in RTTs to 5.3-6.3 ms but adequate service performance, proving the robustness and fault-tolerance of the proposed network architecture.

Table 7: IGW failover and rerouting

Failed IGW	Affected Region	Reroute Path	New RTT (ms)	Service Impact
Cyprus–Tartus	Coastal	Coastal → Homs → Aleppo → Turkey	1.9 + 2.3 + 1.6 = 5.8 ms	Acceptable for non-real-time
Jordan–Damascus	South	Damascus → Homs → Coastal → Cyprus	2.0 + 1.9 + 2.4 = 6.3 ms	Acceptable

Turkey–Aleppo	North	Aleppo → Homs → Damascus → Jordan	2.3 + 2.0 + 1.2 = 5.5 ms	Acceptable

All failover paths remain under 7 ms, which is acceptable for most applications, except ultra-low latency (e.g., remote surgery would need local backup).

4.6. Total Cost of Ownership (TCO)

Table 8 compares the 5-year total cost of ownership (TCO) of a cloud-only architecture and a hybrid edge–cloud architecture, including capital expenditure (CAPEX), operational expenditure (OPEX), and deployment expenses, for a 10 Tbps network infrastructure. Table 8 indicates that the hybrid design significantly reduces the 5-year TCO from USD 7.04 million to USD 3.75 million, primarily by minimizing operational costs associated with bandwidth, despite the higher investment required for edge nodes.

Table 8: 5-year TCO (USD) – cloud-only vs Hybrid (10 Tbps Links, 10.74 Gbps Raw Data)

Cost component	Cloud-only	Hybrid (Edges + Cloud)	Notes
CAPEX			
Homs cloud facility & hardware	$5,000,000	$1,500,000	Smaller cloud in hybrid
Edge nodes (3×)	$0	$450,000	$150k per edge (10 Tbps-capable)
IGW termination equipment (10 Tbps)	$1,000,000	$1,000,000	Higher cost for 10 Tbps
OPEX (annual)	-	-	-
Domestic bandwidth (10 Tbps)	$0.1 × 10,740 × 12 = $12,888	$0.1 × 107.4 × 12 = $129	Per year
International transit (10 Tbps)	$0.2 × 10,740 × 12 = $25,776	$0.2 × 107.4 × 12 = $258	Per year
Maintenance & staffing	$120,000	$80,000	Per year
Electricity (grid + solar)	$50,000	$80,000	Higher due to more powerful edges
5-year OPEX total	($12,888+$25,776+$120,000+$50,000)×5 = $1,043,320	($129+$258+$80,000+$80,000)×5 = $801,935	-
Total TCO (5 years)	$5,000,000+$1,000,000+$1,043,320 = $7,043,320	$1,500,000+$450,000+$1,000,000+$801,935 = $3,751,935	-

4.7. Cost Savings

Hybrid architecture reduces TCO by 47% ($3.29 million saved over 5 years). Even with 10 Tbps links, the hybrid model remains significantly cheaper.

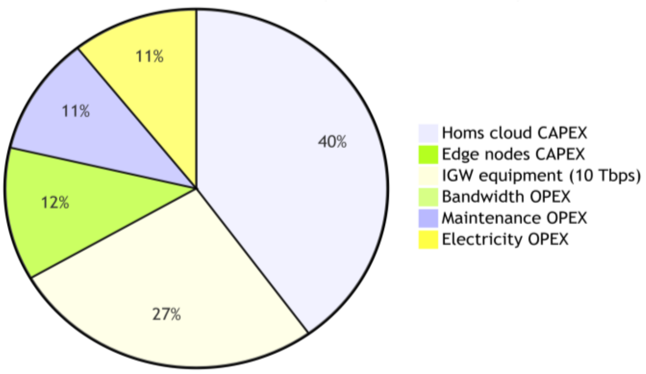

Figure 7: Total cost of ownership distribution for the proposed hybrid edge-cloud architecture with 10 Tbps links; bandwidth costs are negligible (≈0.05%) due to edge preprocessing and low per-Mbps pricing.

If raw data rates increase tenfold (to 107.4 Gbps), cloud-only TCO would skyrocket, while hybrid TCO would scale gracefully. Figure 7 visualizes the TCO breakdown for the hybrid architecture.

4.8. Discussion

Strategic Benefits of 10 Tbps: Switching to a 10 Tbps infrastructure gives the network ample room to grow in the future. Even if data traffic increased by 100 times, the backbone and international gateways would still have ample extra capacity, making them scalable over time. Latency levels remain largely unchanged, while jitter drops significantly due to overprovisioning, improving performance stability and reliability. Also, the presence of 10 Tbps gateways across the country makes it more appealing to hyperscalers, which in turn drives the creation of regional data centers and content delivery networks (CDNs). This infrastructure also strategically positions the network as a potential transit hub between Europe, Asia, and the Middle East, with important geopolitical and economic benefits.

A Comparison with Previous Work: Previous research predominantly examined a maximum capacity of 1 Tbps, thereby constraining their scalability and relevance to forthcoming network requirements. The proposed 10 Tbps architecture, on the other hand, works with new high-capacity submarine cable systems like 2Africa and Blue Raman. This progress shows that, even with ample bandwidth available, edge preprocessing remains very important to prevent wasted network resources by sending raw data over the backbone.

Wrong Things: The proposed architecture has some problems, even though it has some good points. The cost of 10 Tbps network interfaces and optical components remains high, making international gateway equipment more expensive. However, the cost per bit goes down over time. Also, high-capacity switches and routers consume more power, which may lead to underestimated estimates of operational energy costs. There are also regulatory problems, since running high-capacity international links requires appropriate spectrum allocation and cross-border agreements, which this study assumes are in place.

Directions for Future Research: Future research should investigate elastic edge-scaling mechanisms that dynamically allocate additional edge resources when local demand exceeds capacity, despite the unlikelihood of such scenarios in the immediate future. Improvements in multi-band optical switching, especially the use of both the C and L bands, can further boost throughput, reaching 10 Tbps per fiber pair. Also, adding quantum key distribution (QKD) to the high-capacity backbone could make government communications much more secure, opening the door to the next generation of secure networking solutions.

5. Conclusion

The paper has created, modeled, and tested a hyper-distributed edge-cloud architecture that combines three 10 Tbps international gateways and a 10 Tbps domestic backbone. This structure links edge nodes in the Coastal, Damascus, and Aleppo regions to a central cloud in Homs. The system is set up to ensure fast data processing, efficient resource use, and minimal lag between edge and core infrastructure. The proposed system shows significant improvements in latency, with edge nodes achieving a local round-trip time (RTT) of less than 0.5 ms. Using gateways to communicate with people in other countries adds only 1.2 to 2.4 milliseconds of delay. Such low latency enables important real-time applications, such as telemedicine, industrial robotics, and autonomous agriculture. This improves operational efficiency across many areas. The architecture does a good job of optimizing bandwidth by preprocessing traffic at the edge. This cuts backbone traffic, both domestic and international, by 99%, from 10.74 Gbps to 107.4 Mbps. As a result, even high-capacity 10 Tbps links are used effectively without congestion. The system also ensures compliance with all national data protection laws by allowing sensitive data to be stored at edge nodes. Three international gateways also make the system more resilient by providing reliable failover options with round-trip times below 6.5 ms. The hybrid edge-cloud architecture delivers significant cost savings. The total cost of ownership drops by 47% over five years, from $7.04 million to $3.75 million, when compared to a cloud-only model. This is even though edge infrastructure costs more to set up. This shows that the proposed method is both financially sound and capable of growth.

5.1. Recommendations

The Ministry of Communications and Technology should quickly set up edge nodes at the landing points of each international gateway: Tartus (coastal), the Damascus border (Jordan), and the Aleppo border (Turkey). The hardware should be ready for 10 Tbps, as shown in Table 2. The domestic backbone should be upgraded to a 10 Tbps DWDM network that goes through Homs, Damascus, Aleppo, and the coast. To keep latency low, the network should use cut-through switching. All government IoT projects must follow an edge-first data processing policy that limits the sending of sensitive raw data to central or external locations. Also, there should be bilateral peering agreements with Cyprus, Jordan, and Turkey to ensure high-capacity, low-latency connectivity. It is suggested that international partners such as the EU, UNDP, and the World Bank fund a trial run of the coastal edge node and connect it to the Cyprus international gateway for telemedicine applications, such as connecting

Latakia to hospitals in Cyprus. They should also help build capacity by providing Syrian engineers with technical training in Kubernetes, 10 Tbps optical networking, and edge computing. Syria is at a unique crossroads. It can use its existing high-capacity international gateways to move from centralized cloud systems to a distributed, resilient, and independent edge-cloud continuum. This work lays out a strategic plan; the next step is to put it into action in a timely, organized way.

Appendix A: List of Acronyms

- **CAPEX:** Capital Expenditure
- **DWDM:** Dense Wavelength Division Multiplexing
- **FGCS:** Future Generation Computer Systems
- **IGW:** International Gateway
- **IoT:** Internet of Things
- **OPEX:** Operational Expenditure
- **RTT:** Round Trip Time
- **STE:** Syrian Telecommunications Establishment
- **Tbps:** Terabits per second (1000 Gbps)
- **TCO:** Total Cost of Ownership

Appendix A.1: Sensitivity Analysis – Varying International Link Speed (10 Tbps Baseline)

Even at 100 Tbps, transit cost remains negligible (Table A1).

Table A1: TCO sensitivity to international link speed (Hybrid Architecture, 5 Years); Note: Per-Mbps cost decreases with higher speed

International Link Speed	Monthly Cost Per Mbps	5-Years International Transit Cost (Hybrid)	Total TCO (Hybrid)
1 Tbps	\$1.0	\$1.0 × 107.4 × 60 = \$6,444	\$3,758,379
10 Tbps (Baseline)	\$0.2	\$0.2 × 107.4 × 60 = \$1,289	\$3,751,935
100 Tbps	\$0.05	\$0.05 × 107.4 × 60 = \$322	\$3,750,968

Acknowledgment: The authors sincerely acknowledge the support, guidance, and research facilities provided by Al-Shahbaa Private University, Ukrainian National Forestry University, University of Castilla-La Mancha, Shinawatra University, Land Reform Area Development Bureau, and Scien.co.

Data Availability Statement: The data for this study are available upon request from the corresponding author.

Funding Statement: This manuscript and research paper were prepared without any financial support or funding.

Conflicts of Interest Statement: The authors declare no conflicts of interest related to the publication of this research paper.

Ethics and Consent Statement: This research adheres to ethical guidelines and obtains informed consent from all participants.

References

1. M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Commun. ACM*, vol. 53, no. 4, pp. 50–58, 2010.
2. T. B. H. Al-Thani, "Law No. (13) of 2016 on Protecting Personal Data Privacy," *Compliance and Data Protection (CDP) department*, Doha, Qatar, 2016. [Accessed by 18/04/2025].
3. I. Gupta, "Internet Intermediaries and Trade Mark Rights, by Althaf Marsoof," *International Journal of Law and Information Technology*, vol. 28, no. 2, pp. 188–191, 2020.
4. A. A. Chingumbe and V. Mahabi, "Assessment of the Factors Affecting the National ICT Broadband Backbone (NICTBB) Systems Restoration Time," *Tanzania Journal of Engineering and Technology*, vol. 43, no. 2, pp. 23–45, 2024.
5. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, 2016.
6. TeleGeography, "Submarine cable map: Cyprus–Syria," *TeleGeography*, 2025. [Accessed by 12/06/2026].

7. R. Roman, J. Lopez, and M. Mambo, “Mobile edge computing, Fog et al.: A survey and analysis of security threats and challenges,” *Future Gener. Comput. Syst.*, vol. 78, no. 1, pp. 680–698, 2018.
8. B. Varghese, N. Wang; S. Barbhuiya; P. Kilpatrick, and D. S. Nikolopoulos, “Challenges and opportunities in edge computing,” in *Proc. IEEE International Conference on Smart Cloud (SmartCloud),* New York, United States of America, 2016.
9. M. Satyanarayanan, “The emergence of edge computing,” *Computer*, vol. 50, no. 1, pp. 30–39, 2017.
10. R. S. Madhuranthakam, “Implementing data quality assurance frameworks in distributed data engineering workflows,” *AVE Trends in Intelligent Computing Systems*, vol. 1, no. 4, pp. 241–251, 2024.
11. Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, “A Survey on Mobile Edge Computing: The Communication Perspective,” *IEEE Commun. Surv. Tutor.*, vol. 19, no. 4, pp. 2322–2358, 2017.
12. A. K. Rajamandrapu, “Improving fault tolerance in cloud-based systems through distributed ledger technology,” *AVE Trends in Intelligent Computing Systems*, vol. 2, no. 1, pp. 52–61, 2025.
13. A. Yousefpour, C. Fung, T. Nguyen, K. Kadiyala, F. Jalali, A. Niakanlahiji, J. Kong, and J. P. Jue, “All one needs to know about fog computing and related edge computing paradigms: A complete survey,” *J. Syst. Archit.*, vol. 98, no. 9, pp. 289–330, 2019.