International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

A Hub-and-Spoke Roaming Cloud Architecture for Remote, Austere, andDisconnected Environments

Mr Yati Ramchandra Gharat

CTO, Engineering, Litmus It Services Private Limited

Abstract

MEGH (Mobile Edge Gateway on Hybrid Cloud) is a purpose-built, hub-and-spoke cloud architecture designed to bring enterprise-grade cloud services to remote, austere, and disconnected operational environments — including conflict zones, disaster-affected areas, high-altitude terrain, maritime operations, and remote African or jungle field sites. Developed by Yati Gharat at Litmus IT Services Pvt Ltd, MEGH is powered by the Apache CloudStack platform and addresses a critical gap in global cloud service delivery: the absence of reliable, secure, and autonomous cloud infrastructure in locations where permanent internet connectivity cannot be guaranteed.

The architecture consists of a centralized hub — comprising a minimum of three geographically distributed controllers hosted in separate Availability Zones (AZs) — and spoke nodes that operate as fully autonomous, standalone cloud instances at the edge. These spoke nodes execute all workloads independently and synchronize data, logs, and evidence back to the hub once internet connectivity is restored. The system integrates a comprehensive security framework including data-at-rest and data-in-motion encryption, dual multi-factor authentication (2-MFA), and an embedded AI engine for proactive issue resolution and autonomous health management. This paper presents the architectural design, operational philosophy, security model, AI integration, use cases, and future research directions of the MEGH platform.

1. Introduction

The exponential growth of cloud computing over the past two decades has transformed the way organizations store, process, and share information. However, the benefits of cloud infrastructure remain largely inaccessible to operations conducted in geographically remote or connectivity-deprived environments. Military field operations, humanitarian disaster response, scientific expeditions, and development projects in remote regions frequently require real-time data processing, secure authentication, inter-team collaboration, and evidence management — yet lack the stable internet connectivity that conventional cloud services mandate.

Existing approaches to this problem — such as ruggedized laptops, satellite-linked private networks, or intermittent VPN tunnels — are either too limited in capability, too expensive, too slow, or too complex to deploy and manage under field conditions. There is a pressing need for a cloud platform that is simultaneously autonomous enough to function without a live internet uplink, yet integrated enough to synchronize reliably with a central management infrastructure when connectivity returns.

MEGH addresses this gap. Designed and implemented by Yati Gharat of Litmus IT Services Pvt Ltd, MEGH leverages the Apache CloudStack open-source cloud management platform to deliver a portable, self-contained cloud node — the spoke — that provides all the essential services of a full enterprise cloud. The spoke operates completely independently in isolation but is architecturally tethered to a centralized hub consisting of at least three controllers deployed across separate Availability Zones for high availability and geographic redundancy. Upon re-establishing connectivity to the hub, the spoke node securely uploads accumulated data, logs, access records, and operational evidence.

2. Motivation and Target Use Cases

The motivation behind MEGH stems from a fundamental mismatch between the operational realities of field-based organizations and the assumptions baked into commercial cloud services. MEGH is designed specifically for, but not limited to, the following deployment scenarios:

2.1 Conflict and War Zone Operations

In active conflict zones, communication infrastructure is frequently destroyed, jammed, or unavailable. Military units, peacekeeping forces, NGOs, and journalists operating in such environments require secure data storage, communications, situational awareness tools, and evidence management capabilities. MEGH enables these organizations to deploy a standalone cloud node that functions entirely without external connectivity, with all data securely encrypted and protected by dual MFA, ensuring operational security even in adversarial environments.

2.2 Natural Disaster and Humanitarian Response

Floods, earthquakes, cyclones, and wildfires routinely destroy telecommunications infrastructure, leaving relief workers unable to coordinate via conventional digital systems. MEGH spoke nodes can be rapidly deployed to disaster sites, providing first responders with authentication services, resource coordination platforms, medical record management, and evidence collection. Once external connectivity is restored, all operational data is synchronized to the central hub, enabling post-incident analysis and reporting.

2.3 Remote African and Developing-Region Field Sites

Much of sub-Saharan Africa, the Amazon basin, remote Central Asian territories, and Pacific island groups lack reliable internet infrastructure. Development organizations, mining companies, agricultural research programs, and public health campaigns operating in these regions require digital capabilities for project management, data collection, and reporting. MEGH provides a cost-effective, deployable cloud node that functions without relying on local infrastructure.

2.4 High-Altitude and Mountain Expeditions

Scientific expeditions, geological surveys, high-altitude search-and-rescue operations, and mountaineering support teams face extreme environmental conditions alongside poor or absent connectivity. MEGH spoke hardware is designed for rugged deployment and provides the computing backbone for such teams, enabling real-time sensor data collection, team authentication, and communication relay functions even above the connectivity threshold.

2.5 Maritime and Offshore Operations

Offshore oil platforms, research vessels, and maritime patrol operations frequently face intermittent satellite connectivity. MEGH enables these platforms to maintain full cloud capability between satellite windows, synchronizing operational data during connectivity windows without disrupting ongoing workflows.

3. Architectural Overview

The MEGH platform is organized into two primary tiers: the Hub (Central Controllers) and the Spoke (Roaming Cloud Nodes). This hub-and-spoke model has been purposefully designed to decouple edge autonomy from central governance — allowing spoke nodes to operate independently while maintaining architectural alignment with the hub.

3.1 Hub: Central Controller Cluster

The hub constitutes the command and control layer of the MEGH ecosystem. It consists of a minimum of three controller nodes deployed across geographically distinct Availability Zones (AZs). This multi-AZ deployment ensures that no single point of failure — whether a data center outage, power disruption, or natural disaster — can render the hub unavailable.

Each hub controller is a fully functional cloud management instance built on the Apache CloudStack platform. The hub provides the following core services:

- Identity and Access Management (IAM): Centralized authentication and authorization services, including LDAP/AD integration, SAML federation, and RBAC role management.
- Data Ingestion and Synchronization: Upon spoke re-connection, the hub receives encrypted data packages, validates their integrity, and ingests them into the central data store.
- Policy and Compliance Engine: Organizational security policies, data governance rules, and compliance configurations are centrally defined and pushed to spoke nodes during synchronization.
- Monitoring and Observability: Centralized dashboards for spoke node health, usage telemetry, and audit log aggregation.
- AI Engine: An embedded artificial intelligence layer responsible for anomaly detection, predictive maintenance, automated remediation, and intelligent alerting.
- Disaster Recovery and Backup: Cross-AZ replication ensures that data synchronized from spoke nodes is immediately replicated to at least two other AZ instances.

3.2 Spoke: Roaming Cloud Node

The spoke is the deployable, field-ready cloud node. It is designed to operate as a fully autonomous, standalone cloud instance capable of delivering enterprise-grade services without any connectivity to the hub or to the public internet. Each spoke is built on Apache CloudStack and mirrors the service stack of the hub.

Key characteristics of the spoke node include:

- Full service autonomy: Authentication, authorization, compute, storage, networking, and application hosting all function independently of hub connectivity.
- Store-and-forward data architecture: All data generated during disconnected operation — including user transactions, logs, access records, and application data — is stored securely in the local encrypted store and tagged for synchronization upon hub reconnection.
- Lightweight physical footprint: Spoke hardware is designed for portability, ruggedness, and low power consumption, enabling deployment via vehicle, aircraft, or pack carry.
- Automatic hub discovery: Upon re-establishing internet connectivity, the spoke autonomously identifies available hub controllers, authenticates to them using pre-provisioned certificates, and initiates the synchronization protocol.
- Tamper detection: Physical and logical tamper detection mechanisms ensure that any unauthorized access attempt is recorded and the spoke node enters a locked state pending hub verification.

3.3 Connectivity and Synchronization Protocol

The synchronization between spoke and hub follows a store-and-forward paradigm. When the spoke has no hub connectivity, all operations are performed locally and all generated data is written to an encrypted local queue. Upon hub reconnection, the spoke initiates a secure mutual TLS (mTLS) handshake with the available hub controller, verifies certificate chains, and begins a staged data upload. The synchronization process prioritizes:

Priority	Data Category	Rationale
P1	Security and audit logs	Compliance and forensic integrity
P2	Authentication records and MFA tokens	Identity governance and session continuity
P3	Operational data and evidence	Core mission data continuity
P4	Configuration delta updates	Policy and software consistency

4. Apache CloudStack as the Foundational Platform

MEGH is built upon Apache CloudStack, an open-source Infrastructure-as-a-Service (IaaS) cloud management platform. CloudStack was selected as the foundational layer for MEGH for several critical reasons:

- Open-source and vendor-neutral: CloudStack's Apache Software License 2.0 licensing allows Litmus IT Services to adapt, extend, and deploy the platform without licensing restrictions, a crucial consideration for field deployments in regions with limited vendor support.
- Proven enterprise scalability: CloudStack has been deployed at hyperscale by organizations such as Apple, SAP, and multiple tier-1 telecommunications providers. Its stability, feature maturity, and hypervisor-agnostic design make it suitable for both the hub and the spoke.
- Hypervisor flexibility: CloudStack supports KVM, VMware vSphere, Citrix XenServer, and bare-metal provisioning — allowing MEGH to be deployed on diverse spoke hardware configurations depending on the field environment.
- Zone and Pod architecture: CloudStack's native concept of Zones, Pods, Clusters, and Hosts maps naturally to MEGH's multi-AZ hub design and the spoke's autonomous zone construct.
- API completeness: CloudStack's comprehensive REST API enables tight integration with MEGH's AI engine, synchronization protocol, and external orchestration tooling.
- Networking stack: CloudStack's support for VLAN isolation, SDN integration via Nicira/NSX, flat networking, and VPC-style constructs enables MEGH to implement consistent network security postures across hub and spoke environments.

The MEGH software layer — branded and deployed as MEGH — represents a customized distribution of Apache CloudStack with additional modules developed by Litmus IT Services to address the unique requirements of spoke-mode autonomous operation, hub synchronization, embedded AI operation, and field-grade security hardening.

5. Security Architecture

Security is the foundational design principle of MEGH, not an afterthought. The platform is designed to protect sensitive operational, personal, and evidentiary data across three threat surfaces: physical device compromise, data interception in transit, and unauthorized logical access.

5.1 Data at Rest Encryption

All data stored on MEGH spoke nodes — including virtual machine disk images, object storage, database volumes, configuration files, log stores, and the synchronization queue — is encrypted at rest using AES-256-GCM encryption. Key management employs a hierarchical key structure:

- Master Key: Derived from hardware TPM (Trusted Platform Module) attestation, unique to each spoke node.
- Volume Encryption Keys (VEKs): Per-volume keys encrypted with the Master Key, allowing granular key rotation and revocation without full re-encryption.
- Archive Keys: Separate keys used to encrypt the synchronization queue, derivable only in conjunction with a hub-issued authorization token, ensuring that queued data cannot be read even if the spoke hardware is seized without hub connectivity.

5.2 Data in Motion Encryption

All network communications within the MEGH platform — between spoke services, between spoke and hub, and between users and spoke endpoints — are protected using Transport Layer Security 1.3 (TLS 1.3) with mutual certificate authentication (mTLS). Key elements include:

- Certificate Authority (CA): A private MEGH CA issues all spoke and hub certificates. The CA private key is stored in an HSM (Hardware Security Module) at each hub AZ.
- Certificate Pinning: Spoke nodes pin hub controller certificates, preventing man-in-the-middle attacks during synchronization even if an adversary has compromised a local network routing layer.
- Perfect Forward Secrecy: ECDHE key exchange ensures that captured traffic cannot be decrypted retroactively even if long-term keys are later compromised.

5.3 Dual Multi-Factor Authentication (2-MFA)

MEGH implements a dual-layer MFA framework that substantially elevates authentication assurance compared to single-factor or single-MFA approaches. The 2-MFA model requires users to satisfy two independent authentication factors drawn from two separate MFA categories:

- Layer 1 MFA: Time-based One-Time Password (TOTP) via an authenticator application (e.g., Google Authenticator, Microsoft Authenticator). This provides possession-based authentication tied to a registered device.
- Layer 2 MFA: Hardware token or biometric verification. In field deployments, this may manifest as a FIDO2/WebAuthn hardware key (e.g., YubiKey), biometric fingerprint verification, or a context-bound session token derived from a physically present access card.

The dual-layer approach ensures that neither a compromised password, nor a stolen authenticator device, nor a cloned hardware token alone is sufficient to gain access. This is particularly important in adversarial field environments where personnel may be captured or coerced.

In disconnected mode, spoke authentication relies on locally cached credential verifiers (bcrypt-hashed passwords, locally-stored TOTP seeds in an encrypted vault, and pre-positioned hardware tokens). Upon reconnection, the hub reconciles the local session log against the central IAM store and flags any authentication anomalies for review.

6. Embedded AI Engine

A distinctive feature of the MEGH architecture is the presence of an embedded artificial intelligence engine within both the hub and the spoke. Rather than relying solely on human operators — who may be unavailable, exhausted, or undertrained in a field context — MEGH's AI engine provides autonomous operational intelligence across multiple domains.

6.1 Anomaly Detection and Threat Response

The AI engine continuously monitors system metrics, network traffic patterns, authentication events, and application logs to detect anomalous behaviour. Trained on baseline operational profiles specific to MEGH's typical workloads, the engine can identify potential security incidents — such as brute-force authentication attempts, unusual data exfiltration patterns, or hardware tampering signals — and initiate automated responses including session termination, account lockout, and data vault sealing.

6.2 Predictive Maintenance and Self-Healing

In field environments, system administrators may be unavailable or untrained. The AI engine performs predictive analysis of hardware telemetry (disk SMART data, CPU thermals, memory error rates, network interface statistics) and proactively schedules maintenance actions — such as log rotation, cache flushing, VM migration, or storage rebalancing — before failures occur. For recoverable software faults, the engine executes automated remediation playbooks, reducing or eliminating the need for human intervention.

6.3 Intelligent Issue Resolution

The AI engine maintains an embedded knowledge base of common MEGH operational issues, their diagnostic signatures, and their remediation steps. When an issue is detected, the engine attempts automated resolution according to a tiered escalation model: automated fix, guided operator resolution, or flagged for hub-level expert intervention post-synchronization. This tiered model ensures that even operators with limited technical expertise can maintain spoke functionality under field conditions.

6.4 Synchronization Intelligence

When connectivity to the hub is established, the AI engine manages the synchronization process intelligently — prioritizing critical data, throttling bandwidth usage to avoid overwhelming limited connectivity, detecting and resolving data conflicts between spoke and hub states, and estimating synchronization completion times to guide operational planning.

7. Service Stack and Software Components

The MEGH spoke node provides a comprehensive suite of cloud services sufficient for full enterprise operations in a disconnected environment. The service stack is organized into the following layers:

Service Layer	Components	Function
Infrastructure	Apache CloudStack, KVM/Hypervisor, SDN	Compute, storage, network virtualisation
Identity & Access	LDAP, SAML, RBAC, TOTP, FIDO2	AuthN, AuthZ, 2-MFA
Storage	Object store, block volumes, file shares	Data persistence and media storage
Networking	VLAN, VPC, NAT, VPN, SDN overlay	Secure network segmentation
Application Hosting	Container orchestration, VM templates	Mission application deployment

AI Engine	MEGH-AI (embedded ML runtime)	Autonomous ops and security intelligence
Security	TLS 1.3, AES-256, TPM, HSM, PKI	End-to-end data protection
Sync & Logging	Encrypted queue, syslog, audit trails	Evidence and compliance management

8. High Availability and Resilience Design

High availability in MEGH is addressed at both the hub and spoke layers, with distinct strategies appropriate to each tier's characteristics.

8.1 Hub High Availability

The hub's minimum three-controller topology across separate AZs provides N+1 resilience: the platform continues to operate if any single AZ becomes unavailable. Controllers use a consensus-based leader election protocol (based on the Raft distributed consensus algorithm) to coordinate state. All hub data is replicated synchronously to at least two AZs before being acknowledged, ensuring zero data loss in the event of a single AZ failure. Load balancing across controller nodes distributes spoke synchronization traffic and API requests, preventing any single controller from becoming a bottleneck.

8.2 Spoke Resilience in Disconnected Mode

During disconnected operation, the spoke's resilience depends on its local hardware configuration. MEGH recommends spoke deployments with RAID-1 mirrored storage (or equivalent NVMe redundancy), redundant power supplies where physical form factor allows, and watchdog-monitored service health checks that automatically restart failed services. The AI engine's self-healing capabilities complement hardware redundancy by addressing software-layer failures autonomously.

9. Operational Workflow

A typical MEGH deployment lifecycle proceeds through the following phases:

Phase 1: Provisioning

Spoke nodes are provisioned at a staging facility with connectivity to the hub. During provisioning, the spoke receives its identity certificates, encryption key material, pre-positioned credential stores, initial operating system and service images, AI engine models, and organizational policy configurations. Provisioning is fully automated via the MEGH hub management portal.

Phase 2: Field Deployment

The spoke node is transported to the field location — whether by vehicle, aircraft, or on foot — and powered up. Upon boot, the spoke automatically initializes all configured services and begins operating in standalone mode. Users authenticate using the local 2-MFA framework and access cloud services without any internet dependency.

Phase 3: Autonomous Operation

All operations — compute, storage, application access, authentication, logging — proceed entirely on the spoke. The AI engine monitors system health and security in real time. All generated data is persisted locally in encrypted form and queued for eventual synchronization.

Phase 4: Synchronization

When internet connectivity is established — whether via satellite uplink, mobile data, temporary fixed line access, or helicopter-delivered relay — the spoke automatically detects available hub controllers,

establishes a secure mTLS session, and initiates the staged synchronization protocol. Synchronization continues in the background, allowing field operations to proceed uninterrupted. Upon successful synchronization, the hub acknowledges receipt and the spoke clears its synchronization queue.

Phase 5: Post-Mission Review

Following the spoke's synchronization with the hub, operational data, evidence, access logs, and AI incident reports are available for review on the hub management portal. Security teams can conduct forensic analysis of field authentication events, data access patterns, and AI-detected anomalies. Compliance and evidence integrity is verifiable through the cryptographic chain-of-custody maintained throughout the synchronization pipeline.

10. Comparison with Existing Approaches

Criterion	Traditional Cloud	Edge Computing	VPN + Laptop	MEGH
Full disconnected ops	No	Partial	Limited	**Yes**
Enterprise service stack	Yes	Limited	No	**Yes**
Dual MFA	Varies	Rare	No	**Yes**
AI self-healing	Varies	No	No	**Yes**
Secure sync on reconnect	N/A	Partial	No	**Yes**
Field-deployable hardware	No	Yes	Yes	**Yes**
Multi-AZ hub redundancy	Yes	No	No	**Yes**
Open-source foundation	No	Varies	No	**Yes**

11. Future Research Directions

MEGH represents a first-generation platform in a rapidly evolving field. Several avenues of future research and development have been identified:

- Satellite-native integration: Direct integration with Low Earth Orbit (LEO) satellite constellations (e.g., Starlink, OneWeb, Amazon Kuiper) to enable intelligent connectivity windows and seamless transition between disconnected and connected modes.
- Multi-spoke mesh networking: Enabling spoke nodes in proximity to form an ad-hoc mesh network, allowing spoke-to-spoke data exchange and collaborative operations without requiring hub connectivity.

- AI model federation: Federated learning across spoke nodes to continuously improve the AI engine's anomaly detection and predictive maintenance models from field operational data, without centralizing raw data.
- Post-quantum cryptography: Migration of MEGH's cryptographic primitives to post-quantum algorithms (e.g., CRYSTALS-Kyber, CRYSTALS-Dilithium) in anticipation of quantum computing threats to current asymmetric encryption.
- Ultra-lightweight spoke variants: Development of spoke configurations optimized for extreme portability — including single-board computer (SBC) and ruggedized smartphone-class hardware — for individual operator or small-team deployment.
- Regulatory and compliance frameworks: Formal alignment of the MEGH security architecture with international standards including ISO 27001, NIST SP 800-53, and region-specific data sovereignty regulations.

12. Conclusion

MEGH represents a significant advancement in the delivery of cloud infrastructure services to remote, disconnected, and austere operational environments. By combining the architectural robustness of Apache CloudStack with a purpose-engineered hub-and-spoke synchronization model, comprehensive dual-layer MFA security, and an embedded AI operations engine, MEGH provides organizations with a cloud capability that transcends the limitations of internet-dependent service delivery.

The platform's design philosophy — autonomous at the edge, governed at the centre — reflects a clear-eyed assessment of the connectivity realities faced by organizations operating in conflict zones, disaster areas, remote geographies, and other challenging environments. MEGH does not require these organizations to compromise on security, functionality, or data integrity simply because they cannot guarantee a live internet connection.

Developed by Yati Gharat at Litmus IT Services Pvt Ltd, MEGH is offered as a foundational contribution to the growing field of edge cloud computing and disconnected operations infrastructure. The authors invite collaboration, peer review, and further research to advance the state of the art in this critical domain.

Acknowledgements

The author acknowledges the contributions of the Apache CloudStack open-source community, whose foundational platform made MEGH possible. The author also acknowledges the operational insights provided by field personnel, humanitarian organizations, and security professionals whose practical challenges inspired the design principles of MEGH.

References

1. Apache Software Foundation. (2024). Apache CloudStack Documentation. https://docs.cloudstack.apache.org
2. Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., & Brandic, I. (2009). Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. Future Generation Computer Systems, 25(6), 599-616.
3. Dinh, H. T., Lee, C., Niyato, D., & Wang, P. (2013). A survey of mobile cloud computing: Architecture, applications, and approaches. Wireless Communications and Mobile Computing, 13(18), 1587-1611.

4. Farris, I., Taleb, T., Khettab, Y., & Song, J. (2019). A survey on emerging SDN and NFV security mechanisms for IoT systems. IEEE Communications Surveys & Tutorials, 21(1), 812-837.
5. NIST. (2020). NIST Special Publication 800-207: Zero Trust Architecture. National Institute of Standards and Technology.
6. Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30-39.
7. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637-646.
8. Stallings, W. (2022). Cryptography and Network Security: Principles and Practice (8th ed.). Pearson.
9. Yati Gharat, Litmus IT Services Pvt Ltd. (2025). MEGH Platform Technical Architecture Specification. Internal Technical Document.
10. Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. Journal of Internet Services and Applications, 1(1), 7-18.