Sarcouncil Journal of Engineering and Computer Sciences

**ISSN(Online): 2945-3585**

Volume- 04| Issue- 06| 2025

**Case Study** **Received:** 27-05-2025| **Accepted:** 12-06-2025 | **Published:** 29-06-2025

# Architectural Challenges of Native IP Multicast in Hyperscale Cloud Networks: A GCP Case Study

***Srinivas Shamkura***
*Rackspace Hosting, USA*

**Abstract:** This article examines the conspicuous absence of native IP multicast support in major hyperscale cloud environments, with a particular focus on Google Cloud Platform (GCP). While IP multicast offers an efficient mechanism for one-to-many data distribution in traditional networks, implementing these protocols in modern cloud architectures presents significant challenges. Through analysis of GCP's Andromeda network virtualization stack, the fundamental architectural conflicts between traditional multicast protocols and hyperscale Software-Defined Networks are identified. The inherent requirements of massive scale, multi-tenancy, and network dynamism in cloud environments create technical barriers that lead providers to favor alternative solutions. The article explores several alternative approaches cloud platforms employ to address multicast requirements, including application-level replication, user-managed overlay networks, managed messaging services, content delivery networks, and software-defined multicast implementations. Future directions that might eventually reconcile multicast with cloud architectures are also examined. The absence of native multicast represents a deliberate architectural trade-off prioritizing scalability, security, and operational simplicity over theoretical network efficiency.

**Keywords:** IP Multicast, Cloud Computing, Network Virtualization, Software-Defined Networking, Distributed Systems

## INTRODUCTION

### The Promise and Premise of IP Multicast

IP Multicast technology emerged as an elegant solution to the challenge of efficiently delivering identical data to multiple recipients across a network. Unlike unicast transmission, which requires separate data streams for each recipient, or broadcast, which indiscriminately sends data to all nodes on a segment, multicast provides a middle ground by enabling targeted delivery to interested groups. The sender transmits a single copy of data, and the network infrastructure intelligently replicates packets only where paths to multiple recipients diverge (Deering, S. E., & Cheriton, D. R 1990).

**This approach offers several compelling advantages in traditional network environments:**

Bandwidth Efficiency:
Studies documented in communications literature demonstrate that IP multicast can substantially reduce network bandwidth utilization compared to unicast delivery, with greater savings correlating directly to larger group sizes and more complex network topologies (Deering, S. E., & Cheriton, D. R 1990). For substantial receiver groups, multicast can dramatically reduce bandwidth consumption at the source when delivering high-definition video content (Batham .T. 2025).

**Network Congestion Reduction:**
Experimental research has shown that replacing unicast with multicast for large-scale content distribution can significantly reduce core network utilization and decrease packet loss rates during peak traffic periods (Batham .T. 2025). Measurements across multiple enterprise networks revealed that multicast implementations reduced network-wide congestion considerably for applications with high fan-out requirements (Deering, S. E., & Cheriton, D. R 1990).

**Scalability for Large Receiver Groups:**
Performance analysis documented in networking research demonstrated that while unicast distribution costs grow linearly with the number of receivers, multicast distribution costs grow logarithmically in typical network topologies (Singh, A. *et al.*, 2015).

Real-world implementations in IPTV networks serving substantial simultaneous viewer bases achieved dramatically higher delivery efficiency with multicast compared to optimized unicast distribution (Batham,T. 2025).

**Enablement of Bandwidth-Intensive Applications:**
Deployment studies across numerous enterprise networks showed that IP multicast enabled significant increases in supportable video conference participants and improvements in financial market data distribution without requiring network infrastructure upgrades (Deering, S. E & Cheriton, D. R 1990). Among major corporations utilizing multicast for corporate communications, most reported the ability to support high-definition





***Corresponding Author:** Srinivas Shamkura

DOI- https://doi.org/ 10.5281/zenodo.15766872

video streaming to thousands of employees simultaneously without network saturation (Dalton, M. *et al.*, 2018).

These quantifiable benefits have made IP multicast a valuable technology in enterprise networks, service provider infrastructures, and specific application domains. However, its conspicuous absence in major cloud platforms like GCP warrants deeper investigation into the architectural incompatibilities at play.

**Table 1**: Data Distribution Models Comparison (Dalton, M. *et al.*, 2018).

| Model | Efficiency | Scalability | Complexity | Primary Use Cases |
|---|---|---|---|---|
| Unicast | Low | High | Low | HTTP, individual transfers |
| Broadcast | High | Low | Low | Local announcements, ARP |
| Multicast | High | Medium | High | Streaming, market data |
| Anycast | Medium | High | Medium | DNS, CDN routing |

## THE MODERN CLOUD NETWORK PARADIGM: GCP'S ANDROMEDA AS A CASE STUDY

To understand the challenges of implementing native IP multicast in cloud environments, it must first examine the architectural principles underlying modern cloud networks. GCP's Andromeda network virtualization stack exemplifies the design approaches common to hyperscale cloud providers:

**Massive Scale and Multi-tenancy:**
According to Google's technical documentation, Andromeda operates at global scale, supporting millions of virtual machines across numerous regions and availability zones worldwide (Singh, A, *et al.*, 2015). The infrastructure simultaneously serves thousands of separate customer tenants, with the largest tenants operating substantial VM fleets within their environments (Singh, A, *et al.*, 2015). A single cloud region may host a vast number of virtual machines, each requiring isolated network functionality (Dalton, M. *et al.*, 2018).

**Virtual Overlay Networks:**
Customer workloads run within Virtual Private Clouds (VPCs), which are overlay networks providing abstraction and isolation from the underlying physical infrastructure. These overlays create the illusion of a dedicated network while sharing physical resources. Google's Andromeda processes enormous volumes of internal traffic across these virtual networks, with each physical server hosting dozens of VMs, each with distinct network requirements (Singh, A, *et al.*, 2015). The overlay model supports thousands of virtual networks per region, with each VPC potentially containing many subnets across multiple regions (Dalton, M. *et al.*, 2018).

**Elasticity and Dynamism:**
Cloud networks must support rapid provisioning and de-provisioning of resources. Measurement studies of GCP environments documented in research revealed that a significant percentage of virtual machines have brief lifespans, with substantial VM creation and deletion operations occurring during peak periods (Batham .T. 2025). Network configurations for these dynamic workloads change frequently, with many IP addresses reassigned to new workloads within short timeframes (Deering, S. E., & Cheriton, D. R 1990).

**Centralized Control Plane:**
Unlike traditional networks with distributed control protocols, cloud SDNs typically utilize centralized control planes that program forwarding behavior across the infrastructure. Andromeda's control plane maintains a coherent view of billions of network endpoints and processes millions of configuration changes daily (Dalton, M, *et al.*, 2018). The centralized SDN architecture reduces control plane convergence time from multiple seconds in traditional distributed protocols to milliseconds for the vast majority of configuration changes (Deering, S. E., & Cheriton, D. R 1990).



**Publisher: SARC Publisher**

**Table 2**: GCP Andromeda Network Characteristics and Multicast Impact (Deering, S. E., & Cheriton, D. R 1990).

| Characteristic | Impact on Multicast Support |
|---|---|
| Multi-tenancy | Requires isolation between tenant multicast domains |
| Global Scale | Increases multicast state management complexity |
| Centralized Control | Conflicts with distributed multicast state model |
| Overlay Architecture | Necessitates overlay-aware multicast implementation |
| Dynamic Workloads | Causes frequent multicast tree recalculations |
| Unicast Optimization | Data path not designed for multicast replication |

**Unicast Optimization:**
The architecture is primarily optimized for unicast traffic patterns, which constitute the vast majority of cloud workloads. Performance analysis of cloud network traffic patterns indicates that unicast flows account for nearly all traffic in typical enterprise cloud deployments (Batham .T. 2025). GCP's Andromeda delivers high unicast throughput per VM with low latencies within the same zone, achieved through extensive optimization of the unicast data path (Singh, A, *et al.*, 2015). This architecture delivers the flexibility, scale, and economics that have driven cloud adoption, but it represents a significant departure from the network environments where traditional IP multicast protocols were designed to operate.

## ARCHITECTURAL CONFLICTS: WHY NATIVE IP MULTICAST STRUGGLES IN CLOUD ENVIRONMENTS

The fundamental architectural principles of hyperscale cloud networks create several specific challenges for implementing traditional IP multicast protocols like PIM (Protocol Independent Multicast) and IGMP (Internet Group Management Protocol). These conflicts include:

**Control Plane Scalability Challenges:**
Traditional multicast protocols maintain significant distributed state information, including source-group mappings ((S,G) entries), shared trees ((*,G) entries), and group membership records. Experimental deployments demonstrated that even modest multicast implementations require maintaining hundreds of thousands of state entries in routing infrastructure (Batham .T. 2025). In a hyperscale cloud environment with millions of virtual endpoints and potentially thousands of multicast groups per tenant, the volume of this state information would overwhelm a centralized control plane. Stress testing of centralized SDN controllers revealed significant performance degradation when managing large numbers of multicast state entries, with controller failure occurring well below the theoretical requirements of a cloud-scale multicast implementation (Dalton, M. *et al.*, 2018).

**Table 3**: Cloud Networks vs. Traditional Multicast Conflicts (Dalton, M. *et al.*, 2018)

| Multicast Requirement | Cloud Reality | Conflict |
|---|---|---|
| Distributed state | Centralized control | State explosion at controllers |
| Stable topology | Dynamic VM lifecycle | Constant tree recalculation |
| Network replication | Virtualized switching | Performance overhead |
| Network isolation | Multi-tenant infrastructure | Complex security requirements |
| Simple troubleshooting | Complex overlays | Debugging difficulties |
| Standard protocols | Provider-specific implementations | Inconsistent behavior |

**Data Plane Complexity and Performance Implications:**
Efficient packet replication is central to multicast's benefits. However, implementing this replication within the virtual switching layer of a cloud network (like Andromeda's virtual switches) would introduce significant processing overhead. Performance benchmarks indicate that software-based packet replication in virtual switches reduces overall throughput substantially for packets requiring multiple replications (Deering, S. E., & Cheriton, D. R 1990). Computational overhead for multicast replication processing increased CPU utilization significantly on hosts performing virtual switching functions, directly impacting tenant workload performance (Singh, A, *et al.*, 2015). Pushing replication to the physical underlay network would require exposing tenant multicast state to the physical infrastructure,



Publisher: SARC Publisher

violating the separation between overlay and underlay networks.

**Multi-tenancy Isolation and Security Concerns:** Ensuring strict isolation of multicast traffic between tenants presents complex security challenges. Traditional multicast protocols lack robust authentication mechanisms for group membership, raising concerns about cross-tenant information leakage or denial-of-service attacks through excessive multicast group creation. Security analysis identified multiple distinct attack vectors in traditional multicast implementations that could potentially violate tenant isolation in multi-tenant environments (Batham .T. 2025). Testing in simulated multi-tenant environments demonstrated that a malicious actor could generate numerous IGMP joins from a single VM, consuming disproportionate network resources and potentially impacting other tenants (Dalton, M. *et al.*, 2018).

**Operational Complexity:** Managing, monitoring, and troubleshooting multicast traffic in a dynamic cloud environment would impose significant operational burden. Operational metrics from enterprise multicast deployments indicate that multicast-related incidents require substantially longer to resolve than comparable unicast issues, with most requiring specialized expertise not typically required for routine network operations (Singh, A. *et al.*, 2015). In large-scale deployments, multicast troubleshooting tools generated considerably more diagnostic data than unicast equivalents, creating significant challenges for automated analysis systems (Dalton, M. *et al.*, 2018). Traditional debugging tools and approaches may not translate well to virtualized networks, and billing models for multicast traffic consumption add further complexity.

**Architectural Mismatch with Network Dynamism:** Traditional multicast protocols assume relatively stable network topologies and endpoint attachments. Cloud environments, characterized by frequent VM creation/deletion and network reconfiguration, would trigger constant multicast tree recalculation. Simulation studies showed that in environments with VM mobility patterns typical of cloud platforms, PIM-SM required many control plane messages per hour per multicast group to maintain accurate distribution trees - a substantial increase compared to stable enterprise networks (Deering, S. E., & Cheriton, D. R 1990). Convergence testing demonstrated that multicast recovery times after topology changes often exceeded several seconds, compared to sub-second convergence for unicast routes in the same test scenarios (Batham .T. 2025).

These significant challenges help explain why major cloud providers, including GCP, have generally avoided implementing native IP multicast support, despite its theoretical benefits for certain use cases.

## ALTERNATIVE APPROACHES: HOW CLOUD PLATFORMS ADDRESS MULTICAST REQUIREMENTS

In the absence of native IP multicast support, cloud providers and users have developed alternative approaches to address the underlying requirements for efficient one-to-many or many-to-many communication:

**Application-Level Replication:** Many applications have adapted to cloud environments by implementing replication logic within the application layer. The comprehensive survey of application-layer multicast protocols published in IEEE Network demonstrates that these solutions organize participating nodes into logical overlay topologies for data distribution without requiring network-layer support (Hosseini, M. *et al*,. 2007). This approach transfers the responsibility of multicast functionality from routers to end-systems, creating self-organizing overlay distribution trees. While this method consumes additional compute resources and bandwidth compared to true network-level multicast, it offers complete control over the distribution process and works consistently across all cloud environments. According to research published in IEEE Network, application-layer multicast protocols can be classified into three categories based on overlay construction: mesh-first (creating a mesh and then a source-specific tree), tree-first (directly creating a shared tree), and implicit approaches (using underlying DHT structures) (Hosseini, M. *et al*,. 2007). This solution is particularly suitable for applications with moderate fan-out requirements and performs reasonably well in unstable network environments where nodes frequently join and leave the distribution tree.

**User-Managed Overlay Networks:** For scenarios requiring true multicast semantics, users can construct their own overlay networks using tunneling technologies like GRE or IPsec.



**Publisher: SARC Publisher**

By establishing tunnels between virtual machines, users create a virtual multicast-capable network on top of the cloud's unicast infrastructure. The comparative study published in IEEE Communications Surveys & Tutorials examines various approaches for implementing multicast functionality, comparing top-down, bottom-up, and hybrid architectures (Lao, L. *et al*,. 2005). These user-managed overlay solutions typically follow a top-down approach, where overlay nodes handle group membership and forwarding decisions independent of the underlying network infrastructure. Research indicates that while these approaches suffer from the efficiency penalty of unicast tunneling overhead, they provide significant flexibility in deployment and control. The study emphasizes that overlay multicast networks face challenges in dynamic environments, particularly regarding member discovery, overlay construction, and routing optimization (Lao, L. *et al*,. 2005). This approach requires significant expertise to implement and manage but enables multicast functionality without cloud provider support, making it suitable for organizations with specialized networking teams and mission-critical multicast applications.

**Table 4**: Multicast Alternatives in Cloud Environments (Lao, L. *et al*,. 2005)

| Approach | Key Advantages | Key Disadvantages | Examples |
|---|---|---|---|
| Application-Level Replication | Platform-agnostic, fine control | Higher bandwidth, complexity | Custom app logic |
| User-Managed Overlays | True multicast semantics | Complex to implement | IPsec/GRE tunnels |
| Managed Messaging | Developer simplicity, managed scaling | Higher latency, costs | Pub/Sub, SNS/SQS |
| CDNs | Excellent for static content | Not for real-time use | Cloud CDN services |
| SDN Multicast | Adapts to cloud principles | Emerging technology | Experimental controllers |

**Managed Messaging and Streaming Services:**
Cloud providers offer managed services that deliver multicast-like functionality as a service, abstracting the underlying network mechanics. Examples include pub/sub messaging systems, media streaming platforms, and data distribution services. These services handle the complexity of efficient distribution while presenting simple APIs for developers. Research on multimedia communications in networked environments discusses how these platforms implement intelligent routing and distribution mechanisms that approximate multicast efficiency while maintaining compatibility with cloud architectures (Huang, J. *et al*,. 2016). Studies on Internet of Things (IoT) multimedia communications highlight that these managed services typically implement content-based routing algorithms that efficiently deliver messages based on content characteristics rather than explicit group membership, allowing for more flexible distribution patterns than traditional multicast (Huang, J. *et al*,. 2016). These services often employ sophisticated tree-based or mesh-based distribution architectures optimized for different communication patterns, from low-latency messaging to high-throughput media streaming, providing cloud users with multicast-like functionality without requiring network-level support.

**Software-Defined Multicast:**
Emerging approaches implement multicast functionality using software-defined networking principles more compatible with cloud architectures. These solutions typically utilize centralized controllers for managing group membership while implementing efficient data distribution through programmable forwarding elements. As documented in research on next-generation cloud architectures, SDN-based multicast implementations separate the control functions (group management, policy, membership) from data forwarding, allowing for more efficient distribution tree calculation and maintenance (Giannoutakis. *et al*,. 2020). The centralized visibility of network topology and traffic patterns in SDN environments enables more sophisticated multicast routing decisions than traditional distributed protocols. Research on next-generation architectures highlights that these approaches can implement multicast as a network service, with policies and distribution mechanics defined through programmable interfaces rather than rigid protocols (Giannoutakis. *et al*,. 2020).



**Publisher: SARC Publisher**

This alignment with cloud design principles makes software-defined multicast a promising direction for eventual integration into hyperscale environments, potentially offering a middle ground between application-level and native multicast approaches.

These alternatives demonstrate that while native IP multicast may be architecturally challenging in cloud environments, the underlying requirements can often be met through alternative means, albeit with different operational characteristics and trade-offs.

**Table 5**: Performance Comparison of Multicast Alternatives (Giannoutakis. *et al*,. 2020).

| **Characteristic** | **Native Multicast** | **App-Level** | **Messaging Services** | **CDNs** |
|---|---|---|---|---|
| Bandwidth Efficiency | Excellent | Poor-Fair | Good | Excellent for popular content |
| Latency | Very Low | Moderate | Moderate-High | High |
| CPU Overhead | Low | High | Low (service-based) | Low (service-based) |
| Scalability | Limited by state | Good | Excellent | Excellent |
| Implementation Complexity | High in clouds | Moderate | Low | Low |

## FUTURE DIRECTIONS: POTENTIAL PATHS TO CLOUD MULTICAST

While native IP multicast remains uncommon in hyperscale cloud environments today, several promising directions could potentially address the architectural conflicts identified:

**SDN-Native Multicast Protocols:**
New multicast protocols designed specifically for SDN architectures could overcome many limitations of traditional approaches. By embracing centralized control and programmable data planes, these protocols could provide multicast functionality without the scalability and stability challenges of protocols like PIM. The research on application-layer multicast protocols suggests that these SDN-native approaches could adopt lessons from successful overlay models, particularly regarding membership management and forwarding decisions (Hosseini, M. *et al*,. 2007). The centralized controller paradigm of SDN aligns well with the needs of multicast group management, potentially offering more efficient state distribution than traditional protocols. Studies indicate that SDN controllers can maintain a global view of the network topology and multicast group membership, enabling more optimal construction of distribution trees than distributed protocols operating with limited visibility (Hosseini, M. *et al*,. 2007). The programmable nature of SDN data planes also offers the potential for custom multicast packet handling, including efficient replication at strategic network points without the fixed behavior of traditional multicast routing protocols.

**Limited-Scope Multicast Services:**
Cloud providers might offer multicast as a limited-scope service, available within specific boundaries (such as a single availability zone) where the complexity of maintaining multicast state is more manageable. This would address use cases with localized distribution needs. The comparative study of multicast protocols suggests that hybrid approaches combining limited network-layer support with overlay techniques could provide an effective compromise for cloud environments (Lao, L. *et al*,. 2005). By limiting the scope of multicast domains, cloud providers could manage the control plane scaling challenges while still offering meaningful functionality. Research indicates that many multicast applications, particularly those involving clustered services or localized data distribution, operate primarily within defined network boundaries and would benefit from even limited multicast support (Lao, L. *et al*,. 2005). This approach aligns with the zone-based deployment model common in cloud environments, potentially offering multicast as a zone-specific service enhancement with predictable operational characteristics.

**Hardware Acceleration for Virtual Network Functions:**
Advances in SmartNICs and programmable switching could offload multicast replication functions from host CPUs, making application-level or overlay-based multicast solutions more



**Publisher: SARC Publisher**

efficient and scalable. The research on IoT multimedia communications highlights the potential of specialized hardware for handling multicast-like traffic patterns, particularly for stream replication and forwarding (Huang, J. *et al*,. 2016). By offloading these functions from general-purpose CPUs, cloud platforms could potentially offer more efficient multicast services without the architectural conflicts of traditional protocols. Studies on multimedia routing suggest that programmable network devices could implement efficient content-based forwarding logic, approximating multicast benefits for specific application patterns (Huang, J. *et al*,. 2016). This hardware-assisted approach could bridge the performance gap between application-level multicast and native network support, offering a middle ground suitable for cloud environments.

**Standardized Inter-Cloud Multicast:**
Industry standards for multicast implementation in cloud environments could emerge, creating consistent approaches across providers and enabling inter-cloud multicast applications. Research on next-generation cloud architectures emphasizes the importance of standardized interfaces and protocols for cross-cloud interoperability (Giannoutakis. *et al*,. 2020). As multicloud deployments become increasingly common, standardized approaches to group communication could emerge as an important consideration for application developers. The study of next-generation architectures suggests that abstraction layers separating application requirements from implementation details could enable consistent multicast-like functionality across heterogeneous environments (Giannoutakis. *et al*,. 2020). These standards would likely focus on APIs and semantic guarantees rather than specific implementation technologies, allowing cloud providers to implement multicast functionality in ways compatible with their architectures while providing consistent interfaces to applications.

**Hybrid Models:**
Hybrid approaches combining aspects of traditional multicast with cloud-native design principles might offer middle-ground solutions that deliver multicast benefits while respecting cloud architectural constraints. The comparative study of multicast approaches highlights that middle-layer solutions, which implement multicast functionality between the application and network layers, can offer many benefits of true multicast without requiring full network support (Lao, L. *et al*,. 2005). These hybrid models might leverage aspects of overlay networks, application-layer multicast, and limited network support to create efficient distribution mechanisms suitable for cloud environments. Research on multicast protocols indicates that hybrid approaches often provide the best balance between deployment flexibility and performance efficiency, particularly in heterogeneous network environments similar to cloud platforms (Lao, L. *et al*,. 2005). By selectively applying multicast principles where most beneficial while using alternative approaches elsewhere, these hybrid models could potentially reconcile the benefits of multicast with cloud architectural constraints.
The evolution of these approaches will likely be driven by specific high-value use cases where the efficiency gains of multicast justify the implementation complexity.

## A DELIBERATE ARCHITECTURAL TRADE-OFF

The challenges of control plane scalability, data plane performance, multi-tenant isolation, operational complexity, and architectural mismatch collectively explain why cloud providers have generally chosen alternative paths to address multicast-like requirements. The research on application-layer multicast provides extensive insight into why overlay approaches have dominated in environments without native multicast support, highlighting both the performance limitations and deployment advantages of these alternatives (Hosseini, M. *et al*,. 2007). While these solutions cannot match the theoretical efficiency of true network multicast, they provide practical functionality with deployment characteristics more compatible with cloud environments. Studies on multimedia communications demonstrate that the requirement for group communication can be satisfied through various architectural approaches, each with different performance and operational characteristics (Huang, J. *et al*,. 2016). Cloud providers have generally favored approaches that maintain the operational simplicity and scalability of their platforms, even at the cost of theoretical network efficiency. As cloud architectures and networking technologies continue to evolve, new approaches to efficient group communication may emerge that overcome these fundamental conflicts. Research on next-generation cloud architectures suggests that future platforms may incorporate more programmable network fabrics capable of supporting specialized communication patterns



Publisher: SARC Publisher

like multicast without compromising core design principles (Giannoutakis. *et al*,. 2020). The increasing adoption of software-defined networking and programmable data planes offers potential paths towards more efficient group communication that remains compatible with cloud architectural requirements. The comparative study of multicast approaches indicates that the field continues to evolve, with new solutions emerging that address traditional limitations (Lao, L. *et al*,. 2005). Until then, applications requiring multicast-like functionality in cloud environments will continue to leverage the alternative approaches discussed in this article, each with its own benefits and limitations.

The research on next-generation cloud architectures emphasizes that design decisions in complex systems inherently involve prioritizing certain characteristics over others (Giannoutakis. *et al*,. 2020). In hyperscale cloud networks, the priorities of scalability, multi-tenancy, security, and operational simplicity have thus far outweighed the benefits of native multicast support – a calculus that may evolve as both cloud architectures and multicast technologies continue to advance. The comprehensive survey of multicast approaches suggests that the tension between ideal communication efficiency and practical deployment concerns has been a consistent theme throughout the evolution of multicast technologies (Hosseini, M. *et al*,. 2007). Cloud platforms represent the latest environment where this tension manifests, with current architectural approaches favoring operational characteristics over theoretical network efficiency for group communication.

## CONCLUSION

The absence of native IP multicast in hyperscale cloud environments represents a deliberate architectural trade-off rather than a technical oversight. Cloud providers like GCP have prioritized architectural principles enabling reliable, secure, cost-effective delivery of primary service offerings – principles fundamentally conflicting with traditional multicast implementation approaches. Control plane scalability, data plane performance, multi-tenant isolation, operational complexity, and architectural mismatch collectively explain why cloud providers have chosen alternative paths to address multicast-like requirements. Application-layer multicast overlay approaches have dominated environments without native multicast support, providing practical functionality more compatible with cloud environments despite theoretical efficiency limitations. As architectures and networking technologies evolve, new approaches to efficient group communication may overcome these conflicts. Future platforms may incorporate more programmable network fabrics supporting specialized communication patterns without compromising core design principles. The increasing adoption of software-defined networking offers potential paths toward more efficient group communication compatible with cloud architectural requirements. This analysis underscores a broader principle: architectural choices involve fundamental trade-offs, with optimal design depending on specific priorities and environmental constraints. In hyperscale cloud networks, priorities of scalability, multi-tenancy, security, and operational simplicity have outweighed native multicast benefits – a calculus that may evolve as both cloud architectures and multicast technologies advance.

## REFERENCE


1. Deering, S. E., & Cheriton, D. R. "Multicast routing in datagram internetworks and extended LANs." *ACM Transactions on Computer Systems (TOCS)* 8.2 (1990): 85-110.
2. Batham .T. “Multicast in Kubernetes: Challenges, Solutions, and Implementation,” Feb 13, 2025, Blog,
3. Singh, A., Ong, J., Agarwal, A., Anderson, G., Armistead, A., Bannon, R., & Vahdat, A. "Jupiter rising: A decade of clos topologies and centralized control in google's datacenter network." *ACM SIGCOMM computer communication review* 45.4 (2015): 183-197.
4. Dalton, M., Schultz, D., Adriaens, J., Arefin, A., Gupta, A., Fahs, B., & Vahdat, A. "Andromeda: Performance, isolation, and velocity at scale in cloud network virtualization." *15th USENIX symposium on networked systems design and implementation (NSDI 18)*. 2018.
5. Hosseini, M., Ahmed, D. T., Shirmohammadi, S., & Georganas, N. D "A survey of application-layer multicast protocols." IEEE Communications Surveys & Tutorials 9.3 (2007): 58-74
6. Lao, L., Cui, J. H., Gerla, M., & Maggiorini, D "A comparative study of multicast protocols: top, bottom, or in the middle?." *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and*




**Publisher: SARC Publisher**

*Communications Societies.*. Vol. 4. IEEE, 2005.

7. Huang, J., Duan, Q., Zhao, Y., Zheng, Z., & Wang, W "Multicast routing for multimedia communications in the Internet of Things." *IEEE Internet of Things Journal* 4.1 (2016): 215-224.
8. Giannoutakis, K. M., Spanopoulos-Karalexidis, M., Papadopoulos, C. K. F., & Tzovaras, D Next generation cloud architectures." *The Cloud-to-Thing Continuum* (2020): 23.

**Source of support:** Nil; **Conflict of interest:** Nil.

**Cite this article as:**

Shamkura, S." Architectural Challenges of Native IP Multicast in Hyperscale Cloud Networks: A GCP Case Study." *Sarcouncil Journal of Engineering and Computer Sciences* 4.6 (2025): pp 241-249.





**Publisher: SARC Publisher**