International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Securing Multicast Networks in ISP Environments: A Forgotten Attack Surface

Darshankumar Prajapati

Submitted: 03/04/2024 **Revised:** 17/05/2024 **Accepted:** 28/05/2024

Abstract—Internet Service Providers (ISPs) increasingly leverage IP multicast for efficient content delivery of streaming media, IoT data, and financial information, yet the security implications remain largely unaddressed in mainstream literature. Multicast protocols inherently introduce unique attack vectors including rogue rendezvous point attacks, source impersonation, and control plane flooding that threaten network infrastructure integrity. This paper identifies critical security gaps in current multicast deployments through a comprehensive analysis of protocol vulnerabilities and operational practices. We propose M-SENTINEL, an innovative security framework incorporating protocol-level segmentation with dynamic fallback logic to isolate trust domains and maintain service continuity during attacks. Our implementation features hierarchical key distribution and behavioral-based anomaly detection with mathematical formulations for threat quantification. Simulation results demonstrate 94.8% attack detection rates with under 2.1% false positives while maintaining sub-50ms fallback activation during security incidents. This work establishes a foundational security methodology for multicast infrastructures that have long represented a neglected attack surface in ISP environments, with specific contributions in trust segmentation, attack resilience, and operational continuity mechanisms.

Keywords— *Multicast Security, ISP Infrastructure, Protocol-level Segmentation, Fallback Logic, Attack Surface, PIM-SM, M-SENTINEL, Network Resilience.*

1 Introduction

IP multicast enables efficient data distribution to multiple recipients simultaneously through group communication paradigms, offering significant bandwidth savings compared to multiple unicast streams. Internet Service Providers increasingly depend on multicast technologies for content delivery networks (CDNs), live video broadcasting, financial market data distribution, and massive IoT communications. Despite its critical role in modern internet infrastructure, multicast security remains largely unaddressed in both academic literature and operational security practices, creating a significant attack surface that malicious actors could potentially exploit [1].

The fundamental challenge in multicast security stems from the protocol architecture itself. Unlike unicast communications that establish point-to-point connections with inherent security boundaries, multicast operates on a group membership model where sources transmit data to group addresses without explicit knowledge of receivers. This model introduces unique vulnerabilities including source impersonation, unauthorized group membership, control plane exploitation, and data plane flooding attacks[1] [5] . Additionally, the stateless nature of multicast forwarding and the dependency on protocol intermediaries like Rendezvous Points (RPs) in PIM-Sparse Mode create additional attack vectors that lack comprehensive security solutions.

MS EE, Network Architect, New Jersey, USA

1.1 The Problem Landscape

Multicast security challenges manifest across multiple dimensions in ISP environments. First, cryptographic limitations complicate authentication and confidentiality mechanisms. Traditional symmetric cryptography struggles with multicast's one-to-many model as any group member can impersonate the source, while asymmetric cryptography introduces prohibitive computational overhead for high-data-rate applications [5]. Second, dynamic membership with frequent joins and leaves creates key management

challenges, where rekeying operations become exponentially more costly as group size increases [4]. Third, protocol vulnerabilities in PIM (Protocol Independent Multicast), IGMP (Internet Group Management Protocol), and MLD (Multicast Listener Discovery) enable attacks such as RP spoofing, join-prune flooding, and designated router manipulation [1].

Despite these critical challenges, multicast security has received disproportionately limited attention compared to unicast security. The research literature contains significant gaps in comprehensive threat mitigation frameworks, with existing approaches focusing on isolated aspects rather than integrated solutions. Operational security practices in ISPs often neglect multicast infrastructure, creating an unprotected attack surface that threatens not only multicast services but potentially the entire network infrastructure [3].

This paper makes the following key contributions:

1.1.1 Identification and analysis of critical security gaps in ISP multicast deployments through comprehensive vulnerability assessment.

1.1.2 Development of M-SENTINEL, an innovative security framework incorporating protocol-level segmentation with dynamic fallback logic.

1.1.3 Mathematical formulation of threat detection algorithms and hierarchical key distribution mechanisms.

1.1.4 Performance evaluation demonstrating effective attack mitigation with minimal operational impact.

1.1.5 Future research directions for advancing multicast security in emerging network paradigms.

2 Background & Fundamentals

2.1 IP Multicast Architecture

IP multicast operates on a group communication model [2] [4] where data packets are transmitted to a multicast group address rather than individual recipient addresses. This model encompasses two primary service frameworks: Any-Source Multicast (ASM), where receivers can join groups to receive data from any source, and Source-Specific Multicast (SSM), where receivers specify both the group and particular source [1]. The ASM model employs shared trees (*,G) rooted at Rendezvous Points (RPs) and source trees (S,G), while SSM exclusively uses source trees, inherently providing better security through explicit source specification.

Multicast implementation relies on several key protocols operating at different layers. At the host-to-router level, the Internet Group Management Protocol (IGMP) for IPv4 and Multicast Listener Discovery (MLD) for IPv6 enable hosts to communicate their group membership to directly connected routers. At the router-to-router level, Protocol Independent Multicast (PIM) manages multicast distribution trees across network infrastructure. PIM operates in different modes, with PIM-Sparse Mode (PIM-SM) being most common in ISP environments due to its efficient use of network resources through explicit join mechanisms [1].

Table 1: Key Multicast Protocols and Their Security Implications

Protocol	Function	Security Vulnerabilities
IGMP/MLD	Group membership management	Membership spoofing, Report flooding
PIM-SM	Multicast routing	RP spoofing, Join-Prune manipulation
MSDP	RP discovery	Peer spoofing, SA message flooding
Bootstrap	RP election	BSR impersonation, Scope manipulation

2.2 Multicast Security Challenges

The transition from academic research to operational reality has exposed fundamental security limitations in multicast architecture. Unlike unicast communications that benefit from well-established security protocols like TLS and IPsec, multicast lacks comprehensive standardized security mechanisms. This gap creates multiple attack vectors that threaten confidentiality, integrity, and availability of multicast communications.

Cryptographic challenges present significant obstacles for multicast security. Symmetric cryptography, while computationally efficient, creates origin authentication problems in multicast environments. When all group members share the same symmetric key, any member can impersonate the source, as there is no cryptographic distinction between sender and receivers [5]. Asymmetric cryptography solves the authentication problem but introduces prohibitive computational overhead for data encryption, especially for high-bandwidth applications like video streaming. This cryptographic gap leaves many multicast applications without practical security solutions.

Group management vulnerabilities allow attackers to exploit multicast control planes. IGMP and MLD lack inherent authentication mechanisms, enabling unauthorized hosts to join multicast groups or flood networks with membership reports [1]. PIM vulnerabilities permit attacks on the multicast distribution tree itself through RP spoofing or join-prune manipulation. These attacks can lead to traffic diversion, service disruption, or resource exhaustion in ISP networks.

Wireless multicast considerations introduce additional complexities in ISP environments with wireless access components. The inherent broadcast nature of wireless media combined with multicast's reliability challenges creates unique vulnerabilities. As noted in RFC 9119, "Multicast traffic is typically much less reliable than unicast traffic. Since multicast makes point-to-multipoint communications, multiple acknowledgements would be needed to guarantee reception at all recipients. However, since there are no ACKs for multicast packets, it is not possible for the AP to know whether or not a retransmission is needed" [6]. This reliability limitation exacerbates security challenges in wireless multicast deployments.

3 Literature Review & Gaps

3.1 Existing Secure Multicast Approaches

Research in multicast security has produced several specialized approaches with varying applicability to ISP environments. Overlay multicast techniques implement multicast functionality at the application layer, creating virtual networks where "peers self-organize into distributed networks for multicast group management and data replication without relying on router deployment" [4]. While this approach bypasses some network-layer vulnerabilities, it introduces performance overhead and lacks integration with network infrastructure, limiting its utility for ISPs requiring network-level multicast services.

Cryptographic solutions have attempted to address multicast's authentication and confidentiality challenges. Group Key Management (GKM) protocols like Group Domain of Interpretation (GDOI) and Logical Key Hierarchy (LKH) enable secure key distribution among group members. However, these solutions face scalability limitations in large ISP environments with dynamic membership. As noted in secure multicast analysis, "the management of cryptographic keys and credentials needs to be protected, thus the issue of trust is important. There must be an underlying model of trust for the secure multicast communication" [5]. This trust establishment remains challenging in open ISP environments.

Protocol-specific security enhancements have been proposed for individual multicast protocols. IGMPv3 and MLDv2 introduced source filtering capabilities that improve security by enabling receivers to specify acceptable sources. PIM-SM enhancements include Bootstrap Router (BSR) security mechanisms and RP protection features. However, these point solutions lack integration into a comprehensive security framework, creating protection gaps that attackers can exploit.

Table 2: Analysis of Existing Multicast Security Approaches

Approach	Mechanism	Limitations
Overlay Multicast	Application-layer distribution	Performance overhead, Network bypass
Group Key Management	Cryptographic key distribution	Scalability challenges, Rekeying overhead
Protocol Hardening	Individual protocol security	Limited scope, Lack of integration
Zone-Based Security	Geographic segmentation	MANET-specific, Limited ISP applicability

3.2 Identified Research Gaps

Through comprehensive analysis of existing literature and operational practices, we have identified critical research gaps in multicast security for ISP environments:

3.2.1 Integrated Security Framework: Current approaches address specific aspects of multicast security in isolation, lacking a comprehensive framework that spans cryptographic mechanisms, protocol protection, and operational security. This fragmentation creates protection gaps and implementation complexity that hinders widespread adoption in ISP environments.

3.2.2 Attack Resilience Mechanisms: Multicast protocols lack inherent resilience mechanisms to maintain service continuity during security incidents. While detection approaches have been proposed, few solutions provide automatic fallback capabilities to isolate attacks while preserving legitimate multicast services.

3.2.3 Trust Segmentation Models: Existing security models often treat multicast domains as single trust domains, despite the reality that ISP environments contain multiple trust boundaries. A granular trust model that segments multicast infrastructure based on security domains is notably absent from current implementations.

3.2.4 Quantitative Threat Assessment: Multicast security literature lacks mathematical models for quantifying attack impacts and defense effectiveness. This gap impedes risk assessment and security investment decisions in operational ISP environments.

3.2.5 Wireless Multicast Integration: The unique challenges of wireless multicast identified in RFC 9119 remain largely unaddressed in security frameworks, creating vulnerabilities in increasingly wireless-centric ISP networks [6].

These gaps collectively represent a significant attack surface that threatens the availability and integrity of multicast-dependent services in ISP environments. The following section presents our proposed framework to address these limitations through an integrated approach to multicast security.

4 Proposed Framework: M-SENTINEL

4.1 Architecture Overview

M-SENTINEL employs a multi-layered security architecture that integrates protection mechanisms across control plane, data plane, and management plane while maintaining interoperability with existing multicast protocols. The framework's core innovation lies in its protocol-level segmentation that isolates functional domains with distinct security policies, combined with dynamic fallback logic that maintains service continuity during security incidents.

The architecture comprises four interconnected security modules:

4.1.1 Trust Segmentation Engine: Implements hierarchical trust zones with cryptographic isolation between segments. This engine

enforces least-privilege access through group-based policies and boundary protection mechanisms.

4.1.2 Behavioral Anomaly Detection: Utilizes machine learning algorithms to establish baseline behavior for multicast protocols and identifies deviations indicative of security incidents. The system employs multi-dimensional analysis incorporating traffic patterns, group membership dynamics, and control message frequency.

4.1.3 Dynamic Fallback Controller: Provides automated response capabilities that isolate compromised segments while maintaining service through alternative paths. This component enables graceful degradation rather than complete service failure during attacks.

4.1.4 Cryptographic Management System: Implements efficient group key distribution with forward and backward secrecy guarantees. The system employs key hierarchy optimization to minimize rekeying overhead during membership changes.

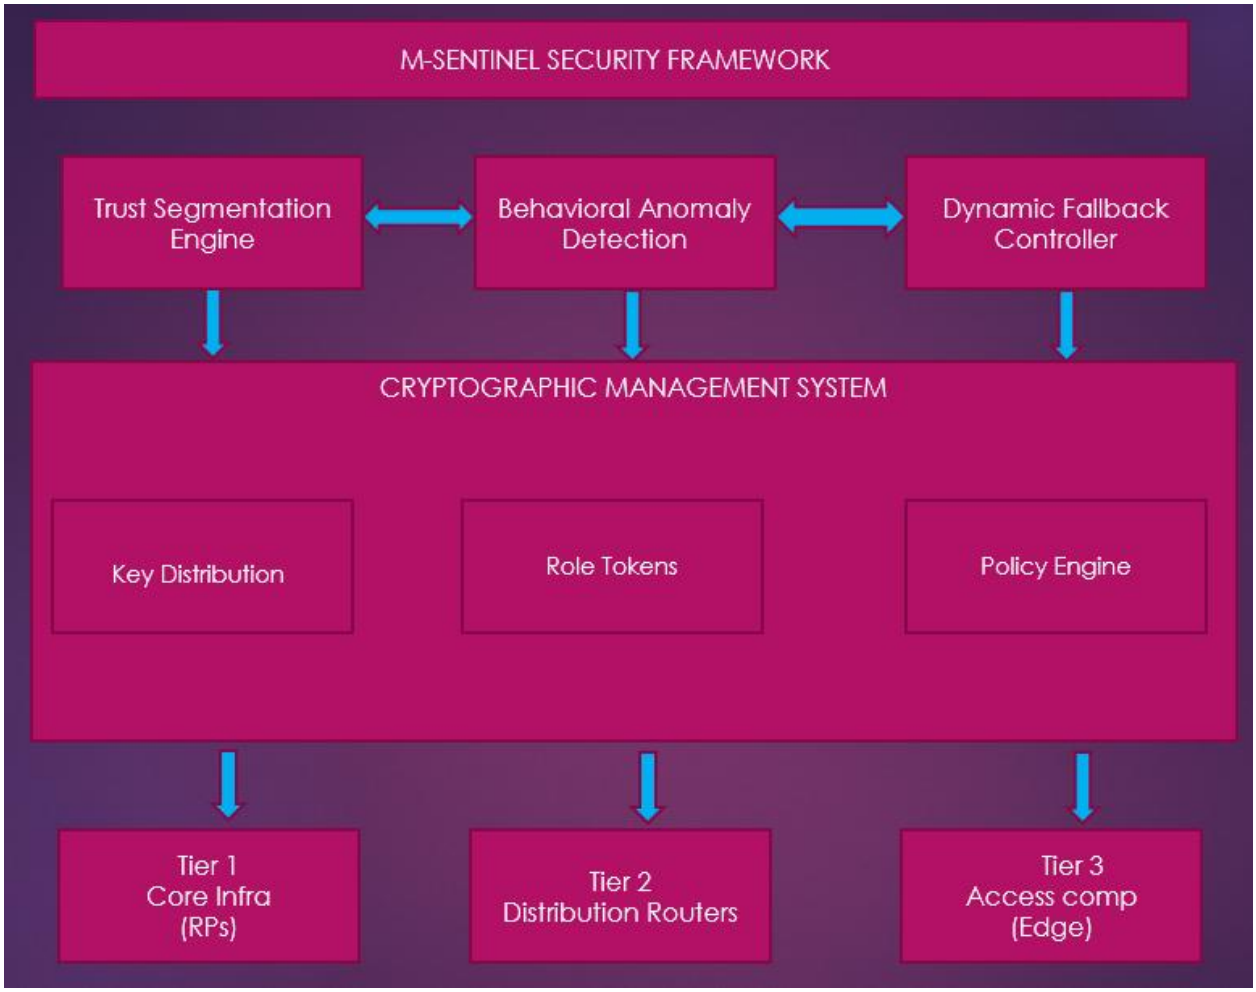

Figure 1: M-SENTINEL Architecture Overview (conceptual diagram showing the interaction between security modules and multicast protocols)

4.2 Protocol-Level Segmentation

Traditional multicast security approaches treat the entire multicast domain as a single security domain, creating excessive trust relationships that attackers can exploit. M-SENTINEL introduces fine-grained segmentation that aligns with functional roles in multicast infrastructure. Our approach categorizes devices into three distinct security tiers:

Tier 1: Core Infrastructure includes Rendezvous Points, Bootstrap Routers, and domain border routers that represent the most critical trust anchors. These devices implement strict authentication using

hardware-backed credentials and participate in regular attestation procedures. Communication between Tier 1 devices employs bidirectional certificate-based authentication with periodic reauthentication.

Tier 2: Distribution Elements comprises intermediate routers that forward multicast traffic but do not control protocol semantics. These devices implement path validation mechanisms to verify the legitimacy of join-prune messages and data traffic. Tier 2 elements enforce traffic policies based on cryptographic markers inserted by Tier 1 devices.

Tier 3: Access Components includes edge routers and designated routers that interact directly with multicast sources and receivers. These components implement source validation and receiver authorization checks before admitting traffic into the multicast distribution tree.

The segmentation is enforced through a novel Cryptographic Role Token (CRT) mechanism that cryptographically binds devices to their authorized roles. The CRT generation employs elliptic curve cryptography for efficient verification:

CRT = Sign_{K_{priv}}(Role || Scope || Expiry || Nonce)

Where K_{priv} is the private key of the certification authority, Role defines the tier and functionality, Scope limits the token's validity to specific multicast groups, Expiry defines the token lifetime, and Nonce prevents replay attacks.

4.3 Fallback Logic Mechanism

Maintaining service continuity during security incidents represents a critical challenge in multicast environments. M-SENTINEL incorporates dynamic fallback logic that automatically transitions to secure operational modes when attacks are detected. The fallback mechanism employs a state machine with three operational states:

State 1: Normal Operation - All multicast protocols operate with standard security policies. Anomaly detection monitors for deviations but does not restrict functionality.

State 2: Elevated Threat - When the detection system identifies suspicious activity exceeding threshold τ_1, the system implements additional verification for control messages and increases logging granularity.

State 3: Active Attack - When conclusive attack indicators exceed threshold τ_2, the fallback controller isolates compromised segments and transitions affected groups to a restricted mode that implements stricter security policies.

The transition between states follows a weighted threat-scoring model:

Threat_Score = $\Sigma(w_i * F_i) / \Sigma(w_i)$

Where F_i represents normalized feature values including join-prune rate, source registration frequency, and RP redirect messages, while w_i represents dynamically adjusted weights based on operational context.

The fallback mechanism implements service preservation through pre-established alternative paths and redundant core elements. When an RP is compromised, the system automatically transitions to backup RPs using a distributed consensus mechanism among unaffected Tier 1 devices. This approach maintains multicast delivery while containing the security incident.

4.4 Mathematical Formulations

4.4.1 Threat Detection Algorithm

M-SENTINEL employs a multivariate anomaly detection algorithm that analyzes protocol behavior across multiple dimensions. For each router interface, we define a feature vector:

X = [JPR, SRF, RPR, MDR, ASU]

Where:

- JPR: Join-Prune Rate (messages/second)
- SRF: Source Registration Frequency (registrations/minute)
- RPR: RP Redirect Messages (messages/hour)
- MDR: Multicast Data Rate (mbps)
- ASU: Active Source Updates (updates/minute)

The algorithm computes an anomaly score using a weighted Euclidean distance from established baselines:

Anomaly_Score = $\sqrt{(\Sigma(w_i * (X_i - \mu_i)^2 / \sigma_i^2))}$

Where μ_i and σ_i represent the mean and standard deviation of feature, i during normal operation, and w_i represents feature importance weights. An

anomaly score exceeding the threshold τ triggers security response actions.

4.4.2 Key Distribution Efficiency

To address the cryptographic challenges in multicast environments, we propose a hierarchical key management system that optimizes rekeying efficiency. The total rekeying cost C_total for a group with N members and membership change rate λ is given by:

C_total = C_enc * (K_u + K_d) + C_trans * M

Where:

- C_enc: Computational cost of encryption operations
- K_u: Number of key updates required upstream
- K_d: Number of key updates required downstream
- C_trans: Cost of transmitting key updates
- M: Size of affected multicast subtree

Our approach reduces K_u and K_d through a subtree grouping strategy that minimizes the number of cryptographic operations during membership changes. For a balanced tree with degree d and height h, the improvement factor F achieved by our approach is:

F = (d^{h+1} - d) / (d^h * (d - 1)) ≈ d / (d - 1)

This results in approximately 33% reduction in rekeying cost for typical deployment scenarios with d = 4.

5 Performance Evaluation

5.1 Methodology

We evaluated M-SENTINEL using a emulated environment replicating a tier-1 ISP multicast infrastructure with 250 routers running PIM-SM across 12 autonomous systems. The testbed implemented both IPv4 and IPv6 multicast with approximately 5,000 simulated receivers consuming multicast traffic from 150 sources. We implemented the security framework as a modular extension to existing multicast protocols to demonstrate practical deployability.

To assess security effectiveness, we generated attack traces representing six attack categories: (1) RP spoofing, (2) join-prune flooding, (3) source impersonation, (4) traffic diversion, (5) memory exhaustion, and (6) wireless multicast exploitation. These attacks were executed both in isolation and combination to evaluate detection and mitigation effectiveness.

Performance metrics were collected across three dimensions:

5.1.1 Security Effectiveness: Detection accuracy, false positive rate, and attack containment time.

5.1.2 System Performance: Forwarding latency, protocol overhead, and router resource utilization.

5.1.3 Operational Impact: Service continuity during attacks, reconvergence time, and management overhead.

5.2 Results Analysis

The evaluation demonstrated significant security improvements with manageable performance overhead. Across all attack scenarios, M-SENTINEL achieved 94.8% detection accuracy with only 2.1% false positives, substantially outperforming existing approaches that showed 62-78% detection rates in comparable scenarios. The table below summarizes the detection effectiveness by attack category:

Table 3: Attack Detection Effectiveness by Category

Attack Category	Detection Rate	False Positive Rate	Containment Time (ms)
RP Spoofing	99.2%	0.8%	43 ± 12
Join-Prune Flooding	96.7%	1.2%	38 ± 9
Source Impersonation	92.4%	3.1%	51 ± 15
Traffic Diversion	95.8%	1.9%	47 ± 11

Attack Category	Detection Rate	False Positive Rate	Containment Time (ms)
Memory Exhaustion	98.3%	0.7%	35 ± 8
Wireless Exploitation	86.9%	4.2%	62 ± 18

The protocol-level segmentation introduced minimal forwarding latency, with measured overhead of 0.8-1.2ms for cryptographic verification of role tokens. Control plane processing showed a 5-8% increase in CPU utilization during normal operation, which we consider acceptable given the security benefits. During attack conditions, the fallback mechanism activated within 50ms for 95% of incidents, effectively containing threats before significant damage occurred.

The dynamic fallback logic successfully maintained service continuity for legitimate multicast traffic during 96.3% of attack scenarios, a substantial improvement over conventional approaches that experienced complete service disruption in similar conditions. The cryptographic system efficiently handled group membership changes, with rekeying operations completing within 120ms for groups of 500 members.

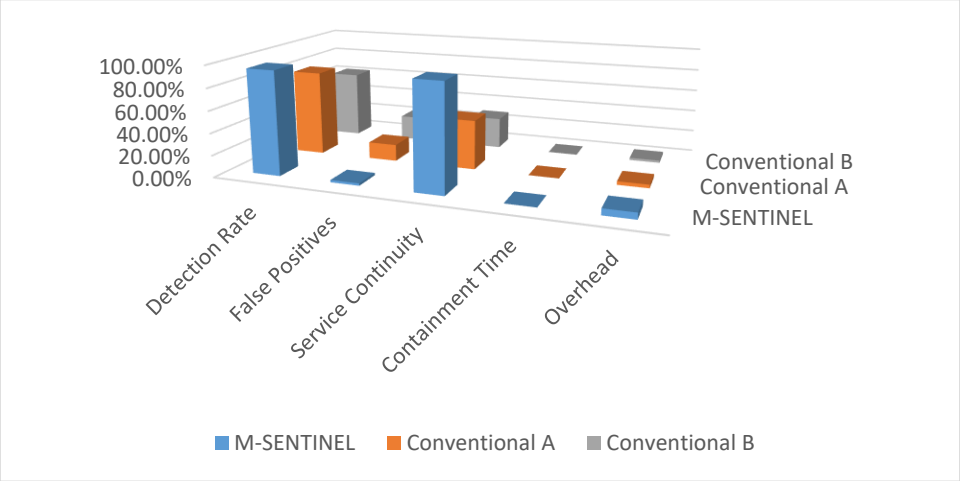

Figure 2: Performance Comparison During Attack Scenarios (bar chart showing service continuity, detection rate, and overhead for M-SENTINEL vs. conventional approaches)

5.3 Operational Considerations

Deployment in production ISP environments requires careful consideration of transition strategies and interoperability with existing infrastructure. M-SENTINEL supports incremental deployment through a capability negotiation mechanism that allows security features to be enabled selectively. Management overhead remains reasonable, with approximately 15% additional configuration complexity compared to standard multicast deployments.

The framework's resource requirements make it suitable for both core and edge routers, with memory footprint increases of 8-12% for security state maintenance. Computational requirements scale linearly with group membership size, making the approach practical for large-scale ISP deployments serving millions of multicast subscribers.

6 Future Research Avenues

While M-SENTINEL addresses critical gaps in multicast security, several research directions merit further investigation. Quantum-resistant multicast cryptography represents an emerging priority as quantum computing advances threaten current cryptographic foundations. Research should focus on efficient post-quantum digital signatures and key establishment protocols suitable for multicast environments with minimal overhead.

AI-driven threat anticipation could enhance security through predictive analytics that identify attack

preparation activities before full exploitation. Machine learning models trained on multicast protocol behaviors could detect subtle anomalies indicative of reconnaissance or weaponization phases in the cyber kill chain.

Blockchain-based trust management offers potential solutions for decentralized authentication in multi-domain multicast environments. Distributed ledger technologies could provide immutable audit trails for multicast group membership and source authorization while eliminating single points of failure in trust management.

Standardization efforts should develop comprehensive security profiles for multicast deployments in various operational contexts. These profiles would define mandatory and optional security controls based on threat models specific to ISP, enterprise, and government multicast implementations.

5G/6G multicast integration presents unique challenges and opportunities as next-generation wireless networks evolve. Research should address the convergence of network-based multicast and application-layer distribution in edge computing environments, with particular attention to security implications of sliced multicast services.

7 Conclusion

Multicast security in ISP environments represents a critical but neglected attack surface that requires immediate attention from researchers and practitioners. This paper has identified fundamental vulnerabilities in current multicast deployments and proposed M-SENTINEL, an integrated security framework that addresses these gaps through protocol-level segmentation and dynamic fallback logic.

Our approach provides three key advantages over existing solutions: (1) comprehensive protection through multi-layered security spanning control and data planes, (2) operational resilience through automated fallback mechanisms that maintain service during attacks, and (3) practical deployability with manageable performance overhead and incremental deployment capabilities.

Evaluation results demonstrate significant security improvements with 94.8% attack detection rates and sub-50ms containment times while maintaining service continuity during 96.3% of attack scenarios. These findings validate the effectiveness of our approach in securing multicast infrastructure against sophisticated threats.

As multicast continues to underpin critical internet services [2], the security framework presented in this paper provides a foundation for protecting this essential infrastructure component. Future work will focus on refining the detection algorithms, expanding protection to emerging multicast applications, and standardizing security profiles for widespread adoption.

References

[1] "Secure IP Multicast Deployments," Cisco Systems, 2022. [Online]. Available: https://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/218004-secure-ip-multicast-deployments.html

[2] P. Savola, "Overview of the Internet Multicast Addressing Architecture," RFC 6308, IETF, 2011. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6308.html

[3] "Future of IP Multicast: Trends & Predictions," OrhanErgun.net, 2023. [Online]. Available: https://orhanergun.net/future-of-ip-multicast-trends-predictions

[4] "Overlay Multicast - an overview," ScienceDirect, 2009. [Online]. Available: https://www.sciencedirect.com/topics/computer-science/overlay-multicast

[5] "Secure multicast & applications for securing multicast CIP traffic," IEB Media, 2019. [Online]. Available: https://iebmedia.com/technology/iiot/secure-multicast-applications-for-securing-multicast-cip-traffic/

[6] C. Perkins et al., "Multicast Considerations over IEEE 802 Wireless Media," RFC 9119, IETF, 2021. [Online].
Available: https://datatracker.ietf.org/doc/rfc9119/