MN IP real-time video Network topology recommendations

1 Overview

Effective deployment of a real-time audio, video, or ancillary data network requires a topology that delivers stable video transmission and dependable control and timing. This white paper details some best practices for video-over-IP infrastructure and defines the operational parameters that ensure Riedel products perform efficiently.

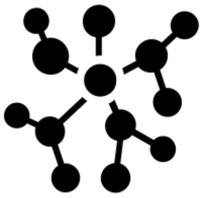

2 Layer Topology Considerations in Real-Time Media Networks

Introduction

In the design of modern real-time audio and video networks, selecting the appropriate network topology is a critical architectural decision that directly affects performance, scalability, and reliability. As production environments evolve—from localized systems to distributed IP-based infrastructures—understanding the implications of operating at Layer 2 versus Layer 3 becomes essential. The choice between these two network layers can have a critical influence on the efficiency of device management, reliability of service discovery, and overall control responsiveness.

Scalability and Discovery Challenges in Layer 2 Media Networks

Layer 2 networks operate at the data link layer, utilizing MAC addresses to forward frames and offering simplicity and low latency. In contrast, Layer 3 networks function at the network layer, employing IP addressing and routing mechanisms that introduce additional complexity but enable greater scalability and flexibility. Note that, Layer 2 networks tend to scale poorly. As the number of connected devices increases, Address Resolution Protocol (ARP) and broadcast traffic can grow exponentially. Each device must receive, process, and respond to ARP requests, which can place significant processing demands on hardware—particularly on devices with limited CPU resources.

Many compact media devices are engineered with high-performance audio and video processing capabilities but are equipped with relatively limited network management processing resources. While such solutions operate flawlessly within a Layer 2 network containing a small number of devices, performance can degrade as the network scales. Internal validations at Riedel have shown that when more than approximately 50 MuoN or FusioN devices are connected within a single Layer 2 domain, some units begin to exhibit slower response times and reduced management performance. This degradation can manifest as delayed command execution and, in severe cases, a complete inability to process control commands. Consequently, when planning a network topology with anticipated growth, it is advisable to adopt a Layer 3 architecture from the outset. In a Layer 3 network, control and management traffic are routed specifically to the relevant devices, ensuring consistent performance regardless of network size.

Moreover, discovery protocols such as Multicast DNS (mDNS) and DNS Service Discovery (DNS-SD) also exhibit limitations when operating within large Layer 2 networks. These mechanisms rely on multicast queries and responses to advertise and discover services, which function efficiently only within small or moderately sized broadcast domains. As the number of devices increases, multicast traffic grows significantly, consuming bandwidth and requiring every device to process each discovery message—even those not intended for it. This additional processing load can affect devices with constrained CPU resources and may lead to slower discovery times or inconsistent visibility of services. By contrast, Layer 3 network architectures allow for the segmentation of multicast domains and the use of controlled discovery relays or gateways, ensuring that mDNS and DNS-SD traffic remain contained and that service discovery remains efficient and predictable at scale.

In Practice: Layer 2 vs. Layer 3

The operational impact of choosing Layer 2 versus Layer 3 becomes particularly clear in practical deployments. Layer 2 networks work well for small, contained systems, but their performance can degrade as device count increases. Layer 3 networks, by contrast, support larger deployments with predictable control and discovery performance. The following table summarizes typical considerations:

Use Cases	Preferred Layer to use	Rationale
Single studio or small facility - under 50 devices	Layer 2	Minimal latency: automatic discovery works efficiently; simple configuration
Campus or multi-room system - under 50 devices	Layer 2 with VLANs	VLANs isolate traffic while maintaining low latency
Large multi-building facility - beyond 50 devices	Layer 3	Routing required for scalability, management, and fault isolation
Remote or cloud production	Layer 3	Supports WAN connectivity, SD-WAN, and VPN; enables centralized management
Broadcast contribution links (ST 2110, AES67)	Layer 3 (with PTP-aware routers)	Standards-based IP transport; scalable timing management; deterministic performance

Operational Observations

Media devices often have limited CPU capacity dedicated to management and control functions. While these devices handle audio and video streams efficiently, excessive Layer 2 broadcast and multicast traffic can overload the management pipeline, particularly in networks with more than 50 devices per broadcast domain. Observed effects include delayed command execution, slower device responses, and, in extreme cases, an inability to process control commands. Layer 3 topologies mitigate these issues by routing management and discovery traffic only to the relevant devices, ensuring predictable operation even as system size increases.

Our recommendations are aligned with the best practices advocated by leading switch manufacturers such as Cisco and Arista.

Conclusion

While Layer 2 networks provide simplicity and low latency for small-scale or isolated deployments, their scalability and management limitations become evident as systems expand. Broadcast and multicast traffic, coupled with limited device processing capacity, can lead to degraded performance in both control and discovery functions. Layer 3 architectures, by contrast, offer structured scalability, improved fault isolation, and better control over multicast behavior, making them more suitable for large or growing real-time media environments.

For new deployments or systems expected to expand over time, a Layer 3 topology should be considered from the design phase onward. Hybrid architectures—combining Layer 2 domains for localized real-time transport and Layer 3 routing for management and control—can provide an effective balance between performance and scalability, ensuring consistent operation even under high system loads.