
Stream: Internet Engineering Task Force (IETF)
RFC: [9706](#)
Category: Informational
Published: December 2024
ISSN: 2070-1721
Authors: L. Giuliano C. Lenart R. Adam
Juniper Networks Verizon GEANT

RFC 9706

TreeDN: Tree-Based Content Delivery Network (CDN) for Live Streaming to Mass Audiences

Abstract

As Internet audience sizes for high-interest live events reach unprecedented levels and bitrates climb to support formats and applications such as 4K, 8K, and Augmented Reality (AR), live streaming can place a unique type of stress upon network resources. TreeDN is a tree-based Content Delivery Network (CDN) architecture designed to address the distinctive scaling challenges of live streaming to mass audiences. TreeDN enables operators to offer Replication-as-a-Service (RaaS) at a fraction of the cost of traditional, unicast-based CDNs -- in some cases, at no additional cost to the infrastructure. In addition to efficiently utilizing network resources to deliver existing multi-destination traffic, this architecture also enables new types of content and use cases that previously were not possible or economically viable using traditional CDN approaches. Finally, TreeDN is a decentralized architecture and a democratizing technology that makes content distribution more accessible to more people by dramatically reducing the costs of replication.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9706>.

Copyright Notice

Copyright (c) 2024 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Requirements Language	3
3. Problem Statement	3
4. Applicability	4
5. Multicast Challenges in the Past	4
6. TreeDN Architecture	5
6.1. TreeDN Overlays	6
6.2. TreeDN Native On-Net	6
7. Replication-as-a-Service (RaaS)	7
8. Decentralization/Democratization of Content Sourcing	8
9. Transport-Layer-Related Differences between TreeDN and Traditional CDNs	8
9.1. Integration with Unicast	8
9.2. Reliability, Adaptive Bitrates, and Congestion Control	8
9.3. Authorization and Encryption	9
10. TreeDN Deployments	9
11. Operational Considerations	10
12. Security Consideration	10
13. IANA Considerations	10
14. References	11
14.1. Normative References	11

14.2. Informative References	11
Appendix A. Netverses	13
Acknowledgements	14
Authors' Addresses	14

1. Introduction

As Internet audience sizes for high-interest live events reach unprecedented levels and bitrates climb to support formats and applications such as 4K, 8K, and Augmented Reality (AR), live streaming can place a unique type of stress upon network resources. TreeDN is a tree-based Content Delivery Network (CDN) architecture designed to address the distinctive scaling challenges of live streaming to mass audiences. TreeDN enables operators to offer Replication-as-a-Service (RaaS) at a fraction of the cost of traditional, unicast-based CDNs -- in some cases, at no additional cost to the infrastructure. In addition to efficiently utilizing network resources to deliver existing multi-destination traffic, this architecture also enables new types of content and use cases that previously were not possible or economically viable using traditional CDN approaches. Finally, TreeDN is a decentralized architecture and a democratizing technology that makes content distribution more accessible to more people by dramatically reducing the costs of replication.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Problem Statement

Live streaming to mass audiences can impose unique demands on network resources. For example, live sporting events that broadcast over the Internet to end users have a much lower tolerance for long playout buffers than typical on-demand video streaming. Viewers of live sporting events have long been conditioned by broadcast television to expect to see the content in real time, with only very short buffers for broadcast delays to prevent profanity and other objectionable content from making on the air (this is known as the "seven-second delay" [BROADCAST-DELAY]). With micro-betting, even this 5 to 10 second delay can be too long. By comparison, when watching on-demand movies, an extra one- or two-minute playout buffer tends to be perfectly acceptable for viewers. If playout buffers for live sports are that long, viewers run the risk of being alerted to a game-winning score from text messages from friends or cheers from the bar across the street minutes before they view it themselves.

Another unique characteristic of live streaming is the join rate. While on-demand video streaming can consume massive amounts of network resources, the viewing rates tend to be smooth and predictable. Service Providers (SPs) observe gradual levels of traffic increases over the evening hours corresponding to prime-time viewing habits. By comparison, viewing rates of live video streams can more closely resemble a step function with much less predictability as mass audiences of viewers tune in to watch the game at the same time.

Previous efforts for more efficient network replication of multi-destination traffic have experienced mixed success in terms of adoption. IP multicast is widely deployed on financial networks, video distribution networks, L3VPN networks, and certain enterprises. However, most of these deployments are restricted to "walled-garden" networks. Multicast over the global Internet has failed to gain traction, as only a very small portion of the Internet is multicast enabled at this time.

TreeDN is a tree-based CDN architecture that is the result of the evolution of network-based replication mechanisms and is based on lessons learned from what has and has not worked well in the past. TreeDN addresses the fundamental issues of what has hindered multicast from adoption on the global Internet and enables SPs the opportunity to deliver new Replication-as-a-Service (RaaS) offerings to content providers, while more efficiently utilizing network resources by eliminating duplicated traffic. Thus, this improves the experience of end users. TreeDN accomplishes this with the combination of a simplified model of native multicast along with network overlays to reach receivers on unicast-only parts of the Internet.

By more efficiently supporting multi-destination traffic, TreeDN is an architecture that can enable new types of content (such as AR live streaming to mass audiences) that previously weren't possible or economically viable on the Internet due to the inefficiencies of unicast.

4. Applicability

While the primary use case mentioned throughout this document is live streaming of multimedia content (e.g., audio, video, AR, and real-time telemetry data), the TreeDN architecture can provide efficient delivery for any content that needs to be replicated and delivered to multiple destinations. For example, large software file updates (e.g., OS upgrades) that need to be delivered to many end users in a very short window of time can cause significant strain on network resources. Using TreeDN, this use case can be handled much more efficiently by the network.

5. Multicast Challenges in the Past

The following issues have been some of the primary challenges for deployment of IP multicast over the global Internet. This is not intended to be an exhaustive list but rather a list that provides context for the solution and how it addresses these primary challenges.

- The "All or Nothing" problem: IP multicast requires every Layer 3 hop between the source and receivers to be multicast enabled. To achieve ubiquitous availability on the global Internet, this essentially means that nearly every interface on every router and firewall

between all end hosts must support a multicast routing protocol (such as Protocol Independent Multicast - Sparse Mode (PIM-SM) [RFC7761] or the Multipoint Label Distribution Protocol (mLDP) [RFC6388]). This requirement creates a bar to deployment that is practically impossible to overcome.

- The "It's Too Complex" problem: Operators have long complained that multicast routing protocols like PIM-SM are simply too complex, making it costly to design, configure, manage, and troubleshoot IP multicast in the network.
- The "Chicken and Egg" problem: There's not much multicast content because there's not much of a multicast-enabled audience, but there's not much of a multicast-enabled audience because there's not much multicast content.

TreeDN is the evolution of network-based replication based on lessons learned over decades and is designed to address the problems listed above.

6. TreeDN Architecture

TreeDN leverages a simplified model for multicast deployment combined with network overlays to deliver traffic to receiving hosts on unicast-only networks. With network overlays, a service can be achieved and delivered to end users while recognizing and tolerating the practical realities of what is possible over a network as diverse as the global Internet. That is, the replication service is available to users and applications across the global Internet regardless of what protocols may exist in the underlying networks that constitute the underlay.

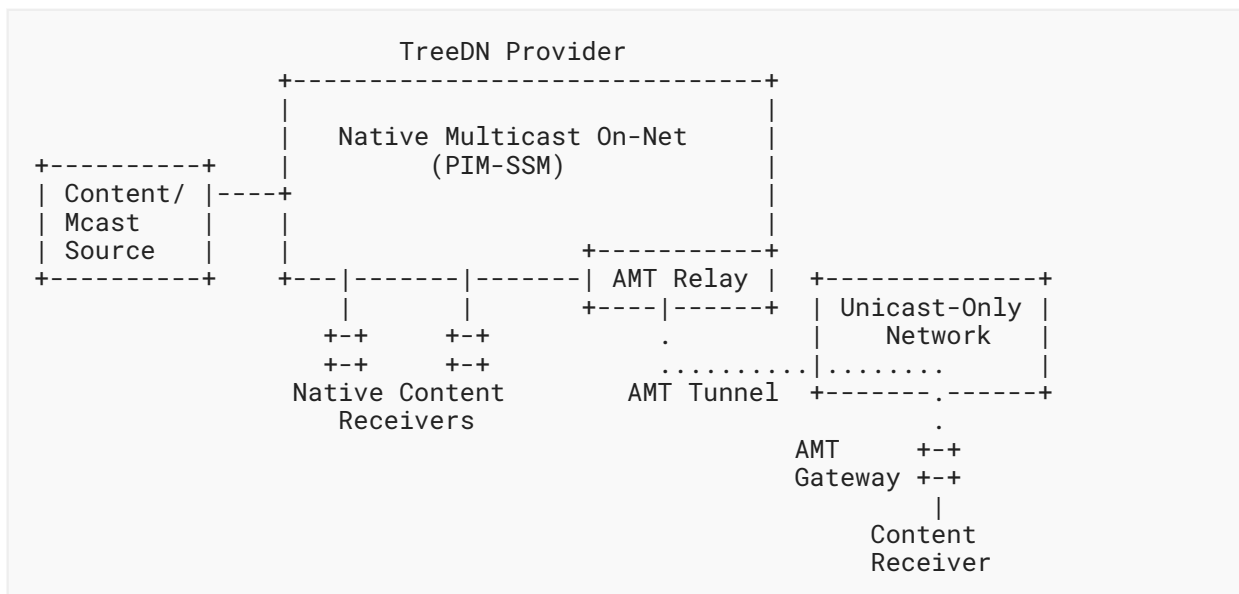


Figure 1: TreeDN Provider Example

6.1. TreeDN Overlays

One overlay technology that TreeDN leverages is Automatic Multicast Tunneling (AMT) [RFC7450]. With AMT, end hosts on unicast-only networks (AMT Gateways) can dynamically build tunnels to routers on the multicast-enabled part of the network (AMT Relays) and receive multicast streams. The AMT Gateway is a thin software client that typically sits on the receiving end host and initiates the tunnel at an AMT Relay. The AMT Relay is a tunnel server that typically sits at the border of the multicast network. AMT allows any end host on the Internet to receive multicast content regardless of whether their local provider supports multicast (aka, "off-net receivers"), which addresses the "All or Nothing" problem. Links and devices that do not support multicast are simply tunneled over -- they no longer present a barrier to the overall replication service for end users. Those networks that do deploy and support multicast, as well as the content providers that serve up multicast content, are able to enjoy the benefits of efficient replication and delivery. Further, these benefits can serve as incentives for operators who do not yet support multicast to enable it on their networks, which is a key benefit of incremental deployment described in Section 4.3 of [RFC9049]. Once the cost of carrying duplicated unicast tunnels is perceived by those operators to exceed the cost of deploying multicast, they are more likely to enable multicast on their networks. Thus, TreeDN effectively supports incremental deployment in a way that was not previously possible with traditional (non-overlay) multicast networking. Finally, AMT also addresses the "Chicken and Egg" problem, as all end hosts on the global Internet that have access to an AMT Relay are capable of becoming audience members.

To support receiving on both native and non-native networks, receiving hosts can first attempt to join the traffic natively, and if no multicast traffic is received, they can fall back to AMT. This fallback mechanism can be handled by the application layer.

In addition to AMT, other overlay technologies like the Locator/ID Separation Protocol (LISP) [RFC9300] can be utilized to deliver content from multicast-enabled networks to end hosts that are separated by portions of the network (at the last/middle/first mile) that do not support multicast.

6.2. TreeDN Native On-Net

Networks that support multicast provide the native on-net component of TreeDN. The primary requirement of the native on-net component is to support Source-Specific Multicast (SSM) [RFC4607]. PIM-SSM, which is merely a subset of PIM-SM, is the multicast routing protocol typically used in SSM. However, any multicast routing protocol capable of supporting SSM can be used in the TreeDN native on-net component, such as mLDP, Global Table Multicast (GTM) [RFC7716], BGP-based Multicast [BGP-MULTICAST], or even BGP Multicast VPN (BGP-MVPN) [RFC6513] for those operators who carry the global routing table in a Virtual Routing and Forwarding (VRF) table. Likewise, any data plane technology that supports SSM, including Bit Index Explicit Replication (BIER) [RFC8279] and Segment Routing (SR) Point-to-Multipoint (P2MP) [RFC9524], can be used.

The key benefit of SSM as the native on-net component of TreeDN is that it radically simplifies the control plane needed to support replication in the network. This simplification comes by moving source discovery from the network layer to some sort of out-of-band mechanism, usually in the application layer. In SSM, the receiver uses the Internet Group Management Protocol Version 3 (IGMPv3) [RFC3376] for IPv4 or the Multicast Listener Discovery Version 2 (MLDv2) protocol [RFC3810] for IPv6 to specify both the source and group address of the multicast stream. This allows the last-hop router to immediately join the multicast stream along the shortest-path tree (SPT) without the need for shared trees. This benefit addresses the "It's Too Complex" problem. By eliminating the need for network-based source discovery, most of the complexity of multicast is then eliminated, which reduces the cost of deploying and operating a multicast network. Further rationale for this SSM-only approach can be found in Any-Source Multicast (ASM) Deprecation [RFC8815].

7. Replication-as-a-Service (RaaS)

Content providers have traditionally used CDNs to distribute content that needs to be delivered to large audiences, essentially outsourcing the task of replication to CDN providers. Most CDNs utilize unicast delivery, as multicast is not an option due to its lack of general availability on the global Internet. TreeDN is a CDN architecture that leverages tree-based replication to more efficiently utilize network resources to deliver simultaneous multi-destination traffic. By leveraging overlay networking to address the "All or Nothing" and "Chicken and Egg" problems, and leveraging SSM to address the "It's Too Complex" problem, TreeDN avoids the practical issues that previously prevented multicast from being a viable option for CDN providers.

TreeDN has several advantages over traditional unicast-based CDN approaches. First, the TreeDN functionality can be delivered entirely by the existing network infrastructure. Specifically, for operators with routers that support AMT natively, multicast traffic can be delivered directly to end users without the need for specialized CDN devices, which typically are servers that need to be racked, powered, cooled, and connected to ports on routers that otherwise could have been consumed by paying customers. In this way, SPs can offer new RaaS functionality to content providers at potentially zero additional cost in new equipment.

Additionally, TreeDN is an open architecture that leverages mature, IETF-specified, and widely implemented network protocols. TreeDN also requires far less coordination between the content provider and the CDN operator. That is, there are no storage requirements for the data, nor group-key management issues, since a TreeDN provider merely forwards packets. A TreeDN provider simply needs to have enough accounting data (e.g., traffic data, number of AMT tunnels, etc.) to properly bill customers for the service. By contrast, traditional unicast-based CDNs often incorporate proprietary, non-interoperable technologies and require significant coordination between the content provider and the CDN to handle such things as file storage, data protection, and key management.

TreeDN introduces a deployment model that requires new considerations for transport-layer mechanisms that are frequently relied upon by traditional unicast-based CDNs. A discussion on these considerations and differences can be found in [Section 9](#).

8. Decentralization/Democratization of Content Sourcing

TreeDN is an inherently decentralized architecture. This reduces the cost for content sourcing, as any host connected to a multicast-enabled network or on a source-capable overlay can send out a single data stream that can be reached by an arbitrarily large audience. By effectively reducing the marginal cost of reaching each additional audience member to zero, from the perspective of the source, TreeDN democratizes content sourcing on the Internet.

9. Transport-Layer-Related Differences between TreeDN and Traditional CDNs

The focus of this document is on the network-layer components that comprise the TreeDN architecture. This section introduces some of the key transport-layer-related differences between TreeDN and traditional unicast-based CDNs that should be taken into consideration when deploying TreeDN-based services. In many cases, these issues are more related to differences between TCP and UDP than differences between unicast and multicast; thus, UDP-based solutions can be leveraged to address most gaps. The aim of this section is to point to some of the existing work to address these gaps, as well as to suggest further work that could be undertaken within the IETF. Further details of these transport-layer mechanisms are beyond the scope of this document.

9.1. Integration with Unicast

Since SSM inherently implies unidirectional traffic flows from one to many, mechanisms that rely on bidirectional communication between receivers and the content provider (such as bespoke advertising, telemetry data from receivers detailing end-user experience, distribution of decryption keys, switching to higher or lower bandwidth streams, etc.) are not well suited to SSM delivery. As such, separate unicast streams between receivers and content providers may be used for this type of "out-of-band" function while SSM is used to deliver the actual content of interest. These "out-of-band" unicast streams **SHOULD** use the same congestion control and authentication mechanisms that are used today for mass audience unicast delivery. Generally speaking, this hybrid unicast-multicast approach is best handled by the application layer and further detail is beyond the scope of this document.

9.2. Reliability, Adaptive Bitrates, and Congestion Control

Traditional unicast-based CDNs frequently rely on HTTPS over TCP transport; thus, they are able to leverage the granularity of TCP-based mechanisms for reliability, congestion control, and adaptive bitrate streaming. However, this granularity comes at a cost of sending a separate data stream to each viewer. Multicast transmissions usually employ UDP, which inherently lacks many of the aforementioned benefits of TCP but can scale much better for mass audiences of simultaneous viewers. Forward Error Correction (FEC) is a mechanism that has demonstrated full recovery for up to 5% packet loss and interruptions up to 400 ms for multicast data streams

in [EUMETSAT-TERRESTRIAL]. NACK-Oriented Reliable Multicast (NORM) [RFC5740] leverages FEC-based repair and other Reliable Multicast Transport (RMT) building blocks to provide end-to-end reliable transport over multicast networks.

QUIC [RFC9000] is another popular transport used by traditional unicast-based CDNs. While QUIC does use UDP, it does not currently support multicast. Multicast extensions to QUIC have been proposed in [QUIC-Multicast].

Section 4.1 of [RFC8085] describes how a sender can distribute data across multiple multicast source-group channels so that each receiver can join the most appropriate channels for its own reception rate capability, thus providing adaptive bitrate capabilities for multicast streams. [DVB-MABR] and [MAUD] extensively describe an architecture that enables reliability and dynamic bitrate adaptation.

TreeDN deployments **MUST** follow the congestion control guidelines described in Section 4.1.4.2 of [RFC7450]. A multicast application that is being distributed over TreeDN deployments **SHOULD** implement congestion control for its data transmission as described in Section 4.1 of [RFC8085]. The AMT gateway **SHOULD** use the topologically closest AMT relay. Section 3.1 of [RFC8777] describes a set of procedures for optimal relay selection.

9.3. Authorization and Encryption

A multicast sender typically has little to no control or visibility about which end hosts may receive the data stream. Encryption can be used to ensure that only authorized receivers are able to access meaningful data. That is, even if unauthorized end hosts (e.g., non-paying end hosts) receive the data stream, without decryption keys, the data is useless. [GKM-IKEv2] describes an extension to the Internet Key Exchange Protocol Version 2 (IKEv2) for the purpose of group key management. [DVB-MABR] and [MAUD] extensively describe an architecture that includes encryption of multicast streams.

10. TreeDN Deployments

EUMETCast Terrestrial is a service from the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) that delivers meteorological satellite data to end users for purposes such as operational monitoring of climates and detection of global climate changes. EUMETCast Terrestrial connects to the GEANT network, which provides TreeDN services to deliver this real-time data natively to end users on multicast-enabled networks and to end users on unicast-only networks via a global deployment of AMT relays. Details of the EUMETCast Terrestrial service over the GEANT TreeDN network are described in [EUMETCast-TERRESTRIAL-AMT]. Additional details on how this deployment uses encryption, authorization, reliability, and unicast feedback channels for end-to-end file delivery monitoring can be found in [EUMETSAT-TERRESTRIAL].

The Multicast Menu [Multicast-Menu] is a web-based portal that can list and launch active multicast streams that are available on a global TreeDN network of various research and education networks. Details of this TreeDN network, as well as the Multicast Menu, are described in [Offnet-Sourcing-Multicast-Menu].

The RARE network is a global testbed interconnecting several National Research and Education Networks (NRENs) via routers running BIER. AMT relays are deployed to deliver multicast traffic from sources on the RARE network to receivers on unicast-only networks across the Internet. Details of the RARE network are described in [[BIER-AMT-Deployment](#)].

11. Operational Considerations

TreeDN is essentially the synthesis of SSM plus overlay networking technologies like AMT. As such, any existing tools to manage, operate, and troubleshoot a PIM-SSM domain and an AMT deployment can be used to manage a TreeDN deployment. Protocol error handling for PIM-SSM can be found in [[RFC4607](#)] and in [Section 4.8](#) of [[RFC7761](#)]; for AMT, it can be found in [[RFC7450](#)].

One potential operational benefit of a multicast-based approach like TreeDN over a traditional, unicast-based CDN is the visibility that multicast state provides in the routing infrastructure. That is, multicast routers maintain a forwarding cache of multicast flows that usually includes the source address, group address, incoming/outgoing interfaces, and forwarding rate. Generally speaking, such flow state information is not typically available in core networks for unicast, so additional tools outside the routing infrastructure are usually required for monitoring CDN performance and troubleshooting issues like packet loss location. Of course, this benefit comes at a cost of additional state being maintained in the routers for multicast.

Additionally, since multicast leverages Reverse Path Forwarding (RPF), the source of the content can potentially have a greater influence over the path taken through the network from source to native receivers/AMT relays. That is, the BGP peer advertising the reachability of the source's subnet can do so in ways where a particular path through the network is preferred for multicast distribution; these methods are not as easy to accomplish with traditional, destination-based unicast routing.

12. Security Consideration

Since TreeDN is essentially the synthesis of SSM plus overlay networking technologies like AMT, the TreeDN architecture introduces no new security threats that are not already documented in SSM and the overlay technologies that comprise it. In particular, [Section 6](#) of [[RFC7450](#)] candidly notes that AMT, like UDP, IGMP, and MLD, provides no mechanisms for ensuring message delivery or integrity, nor does it provide confidentiality, since sources/groups joined through IGMP/MLD could be associated with the particular content being requested.

[[RFC4609](#)] and [[RFC8815](#)] describe the additional security benefits of using SSM instead of ASM.

13. IANA Considerations

This document has no IANA actions.

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <<https://www.rfc-editor.org/info/rfc4607>>.
- [RFC6388] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, DOI 10.17487/RFC6388, November 2011, <<https://www.rfc-editor.org/info/rfc6388>>.
- [RFC7450] Bumgardner, G., "Automatic Multicast Tunneling", RFC 7450, DOI 10.17487/RFC7450, February 2015, <<https://www.rfc-editor.org/info/rfc7450>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

14.2. Informative References

- [Algorhyme] Wikipedia, "Radia Perlman", September 2024, <https://en.wikipedia.org/w/index.php?title=Radia_Perlman&oldid=1245484160>.
- [BGP-MULTICAST] Zhang, Z. J., Giuliano, L., Patel, K., Wijnands, I., Mishra, M. P., and A. Gulko, "BGP Based Multicast", Work in Progress, Internet-Draft, draft-ietf-bess-bgp-multicast-08, 3 June 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-bgp-multicast-08>>.

-
- [BIER-AMT-Deployment]** Mate, C. and F. Loui, "BIER & AMT implementation", IETF 112 Proceedings, November 2021, <<https://datatracker.ietf.org/meeting/112/materials/slides-112-mboned-bier-amt-depoyment-in-geantrare-network-00>>.
- [BROADCAST-DELAY]** Wikipedia, "Broadcast delay", May 2024, <https://en.wikipedia.org/w/index.php?title=Broadcast_delay&oldid=1225656951>.
- [DVB-MABR]** DVB Project, "Adaptive media streaming over IP multicast", DVB Document A176 Rev.3 (Fourth edition), March 2023, <https://dvb.org/wp-content/uploads/2022/01/A176r3_Adaptive-Media-Streaming-over-IP-Multicast_Interim-Draft-TS-103-769-v121_March_2023.pdf>.
- [EUMETCast-TERRESTRIAL-AMT]** Britton, R. and R. Adam, "EUMETCast Terrestrial over AMT", IETF 115 Proceedings, September 2022, <<https://datatracker.ietf.org/meeting/115/materials/slides-115-mboned-eumetcast-over-amt>>.
- [EUMETSAT-TERRESTRIAL]** Espanyol, O., "EUMETSAT Terrestrial Service", IETF 110 Proceedings, February 2021, <<https://datatracker.ietf.org/meeting/110/materials/slides-110-mboned-eumetsat-multicast-over-the-mbone-00>>.
- [GKM-IKEv2]** Smyslov, V. and B. Weis, "Group Key Management using IKEv2", Work in Progress, Internet-Draft, draft-ietf-ipsecme-g-ikev2-19, 28 December 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-g-ikev2-19>>.
- [MAUD]** Nilsson, M. E., Turnbull, R. S., Stevens, T. S., and S. Appleby, "Multicast-Assisted Unicast Delivery", IBC2023 Tech Papers, September 2023, <<https://www.ibt.org/technical-papers/ibt2023-tech-papers-multicast-assisted-unicast-delivery/10235.article>>.
- [Multicast-Menu]** "Multicast Menu", <<https://menu.treedn.net>>.
- [Offnet-Sourcing-Multicast-Menu]** Delwiche, L., "Offnet Sourcing with the Multicast Menu", IETF 114 Proceedings, July 2022, <<https://datatracker.ietf.org/meeting/114/materials/slides-114-mboned-offnet-sourcing-with-the-multicast-menu-01>>.
- [QUIC-Multicast]** Holland, J., Pardue, L., and M. Franke, "Multicast Extension for QUIC", Work in Progress, Internet-Draft, draft-jholland-quic-multicast-05, 7 July 2024, <<https://datatracker.ietf.org/doc/html/draft-jholland-quic-multicast-05>>.
- [RFC4609]** Savola, P., Lehtonen, R., and D. Meyer, "Protocol Independent Multicast - Sparse Mode (PIM-SM) Multicast Routing Security Issues and Enhancements", RFC 4609, DOI 10.17487/RFC4609, October 2006, <<https://www.rfc-editor.org/info/rfc4609>>.
- [RFC5740]** Adamson, B., Bormann, C., Handley, M., and J. Macker, "NACK-Oriented Reliable Multicast (NORM) Transport Protocol", RFC 5740, DOI 10.17487/RFC5740, November 2009, <<https://www.rfc-editor.org/info/rfc5740>>.
- [RFC6513]** Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
-

- [RFC7716] Zhang, J., Giuliano, L., Rosen, E., Ed., Subramanian, K., and D. Pacella, "Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures", RFC 7716, DOI 10.17487/RFC7716, December 2015, <<https://www.rfc-editor.org/info/rfc7716>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8777] Holland, J., "DNS Reverse IP Automatic Multicast Tunneling (AMT) Discovery", RFC 8777, DOI 10.17487/RFC8777, April 2020, <<https://www.rfc-editor.org/info/rfc8777>>.
- [RFC8815] Abrahamsson, M., Chown, T., Giuliano, L., and T. Eckert, "Deprecating Any-Source Multicast (ASM) for Interdomain Multicast", BCP 229, RFC 8815, DOI 10.17487/RFC8815, August 2020, <<https://www.rfc-editor.org/info/rfc8815>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.
- [RFC9049] Dawkins, S., Ed., "Path Aware Networking: Obstacles to Deployment (A Bestiary of Roads Not Taken)", RFC 9049, DOI 10.17487/RFC9049, June 2021, <<https://www.rfc-editor.org/info/rfc9049>>.
- [RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)", RFC 9300, DOI 10.17487/RFC9300, October 2022, <<https://www.rfc-editor.org/info/rfc9300>>.
- [RFC9524] Voyer, D., Ed., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Replication for Multipoint Service Delivery", RFC 9524, DOI 10.17487/RFC9524, February 2024, <<https://www.rfc-editor.org/info/rfc9524>>.
- [Trees] Kilmer, J., "Trees", Poetry Foundation, <<https://www.poetryfoundation.org/poetrymagazine/poems/12744/trees>>.

Appendix A. Netverses

With inspiration from (and apologies to) Radia Perlman [Algorhyme] and Joyce Kilmer [Trees], the following poem is not intended to provide any normative or informative technical value on TreeDN beyond (mild) amusement for the reader who made it this far in the document:

I think that I shall never see
A CDN more lovely than a tree.

A tree whose crucial property
Is efficient mass-audience delivery.

Using SSM for simplified operation
Of native branches that eliminate duplication.

A tree extended by AMT,
Enabling unicast-only receivers full delivery.

A tree that scales to reach millions of places
To viably support the highest of bitrate use cases.

A CDN is built by folks like me,
But only end users can generate enough demand to necessitate a tree.

Acknowledgements

Many thanks to those who have contributed to building and operating the first TreeDN network on the Internet, including Pete Morasca, William Zhang, Lauren Delwiche, Natalie Landsberg, Wayne Brassem, Jake Holland, Andrew Gallo, Casey Russell, Janus Varmarken, Csaba Mate, Frederic Loui, Max Franke, Todor Moskov, Erik Herz, Bradley Cao, Katie Merrill, Karel Hendrych, Haruna Oseni, and Isabelle Xiong. The writing of this document to describe the TreeDN architecture was inspired by a conversation with Dino Farinacci and Mike McBride. Thanks also to Jeff Haas, Vinod Kumar, Ron Bonica, Jeffrey Zhang, and Éric Vyncke for their thoughtful reviews and suggestions, Chris Lemmons for his detailed shepherd review, and Stephen Farrell, Magnus Westerlund, Reese Enghardt, Jurgen Schonwalder, Carlos Pignataro, Erik Kline, Gunter Van de Velde, Warren Kumari, and Zaheduzzaman Sarker for their last call reviews.

Authors' Addresses

Lenny Giuliano

Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
United States of America
Email: lenny@juniper.net

Chris Lenart

Verizon
22001 Loudoun County Parkway
Ashburn, VA 20147
United States of America
Email: chris.lenart@verizon.com

Rich Adam

GEANT

City House

126-130 Hills Road

Cambridge

CB2 1PQ

United Kingdom

Email: richard.adam@geant.org