Multicast Webinar

Thomas Kernen

Agenda

- Multicast vs. unicast delivery
- Multicast enabled deployments
- Addressing, replication and optimisation
- IPv6 & multicast
- Growth of Internet based video services
- Automatic IP Multicast Tunneling

Unicast vs. Multicast

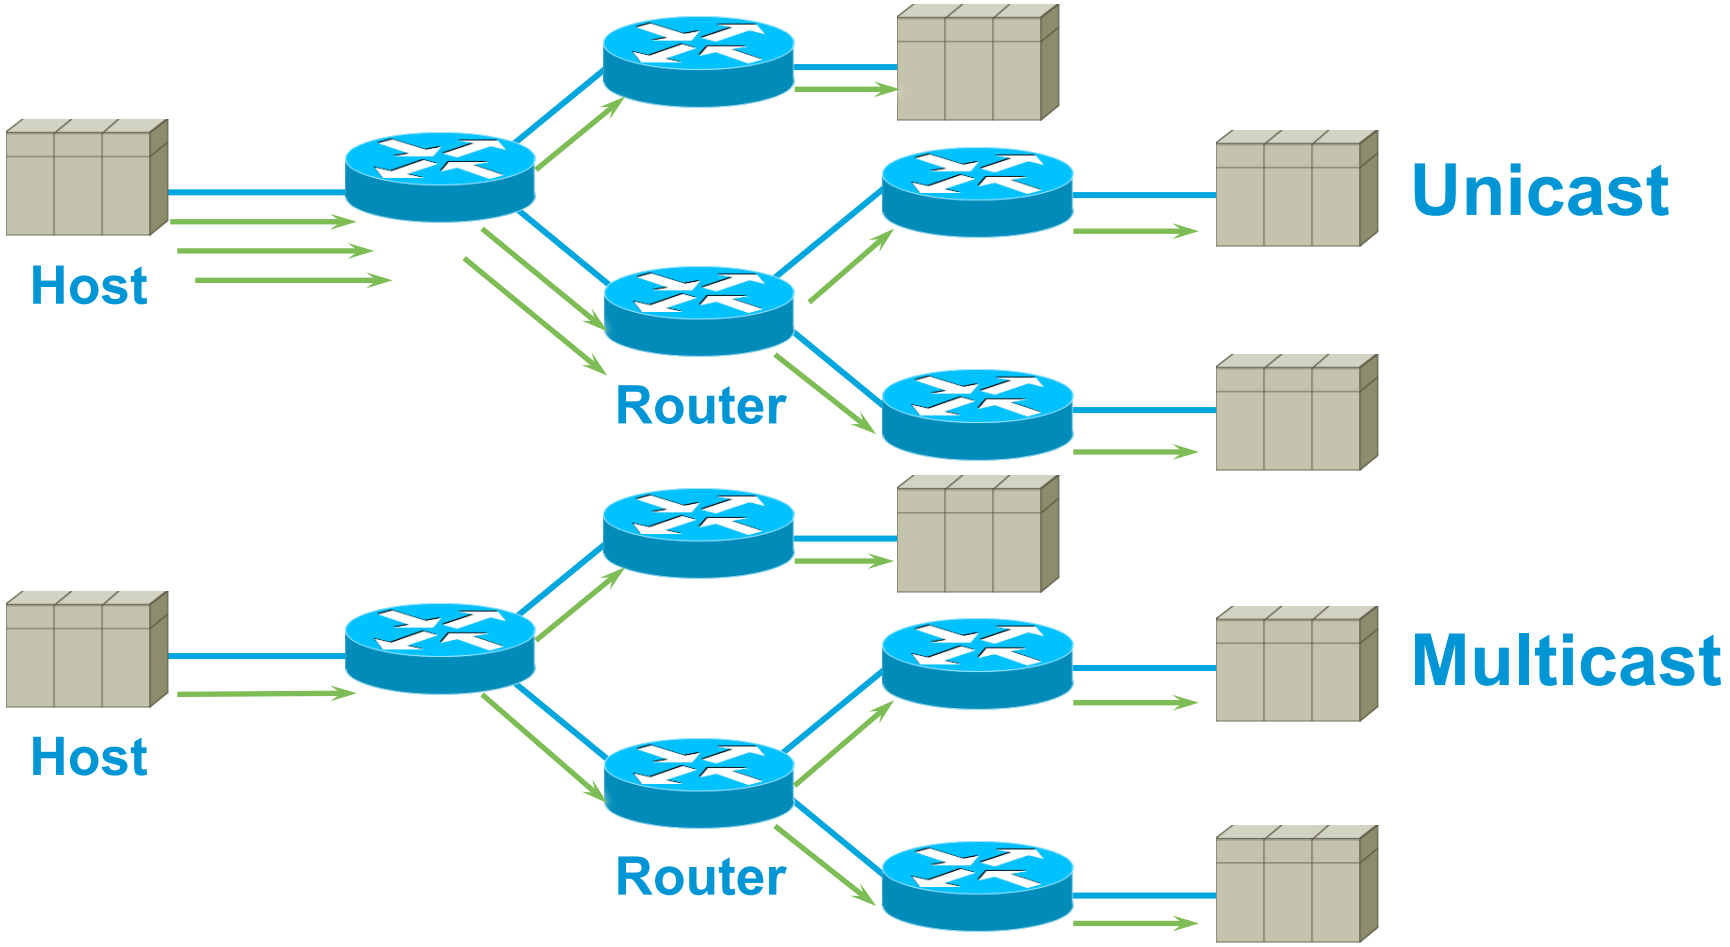

Today we think only of distribution trees when thinking of multicast.
But this is not how it started…

Multicast Application Types

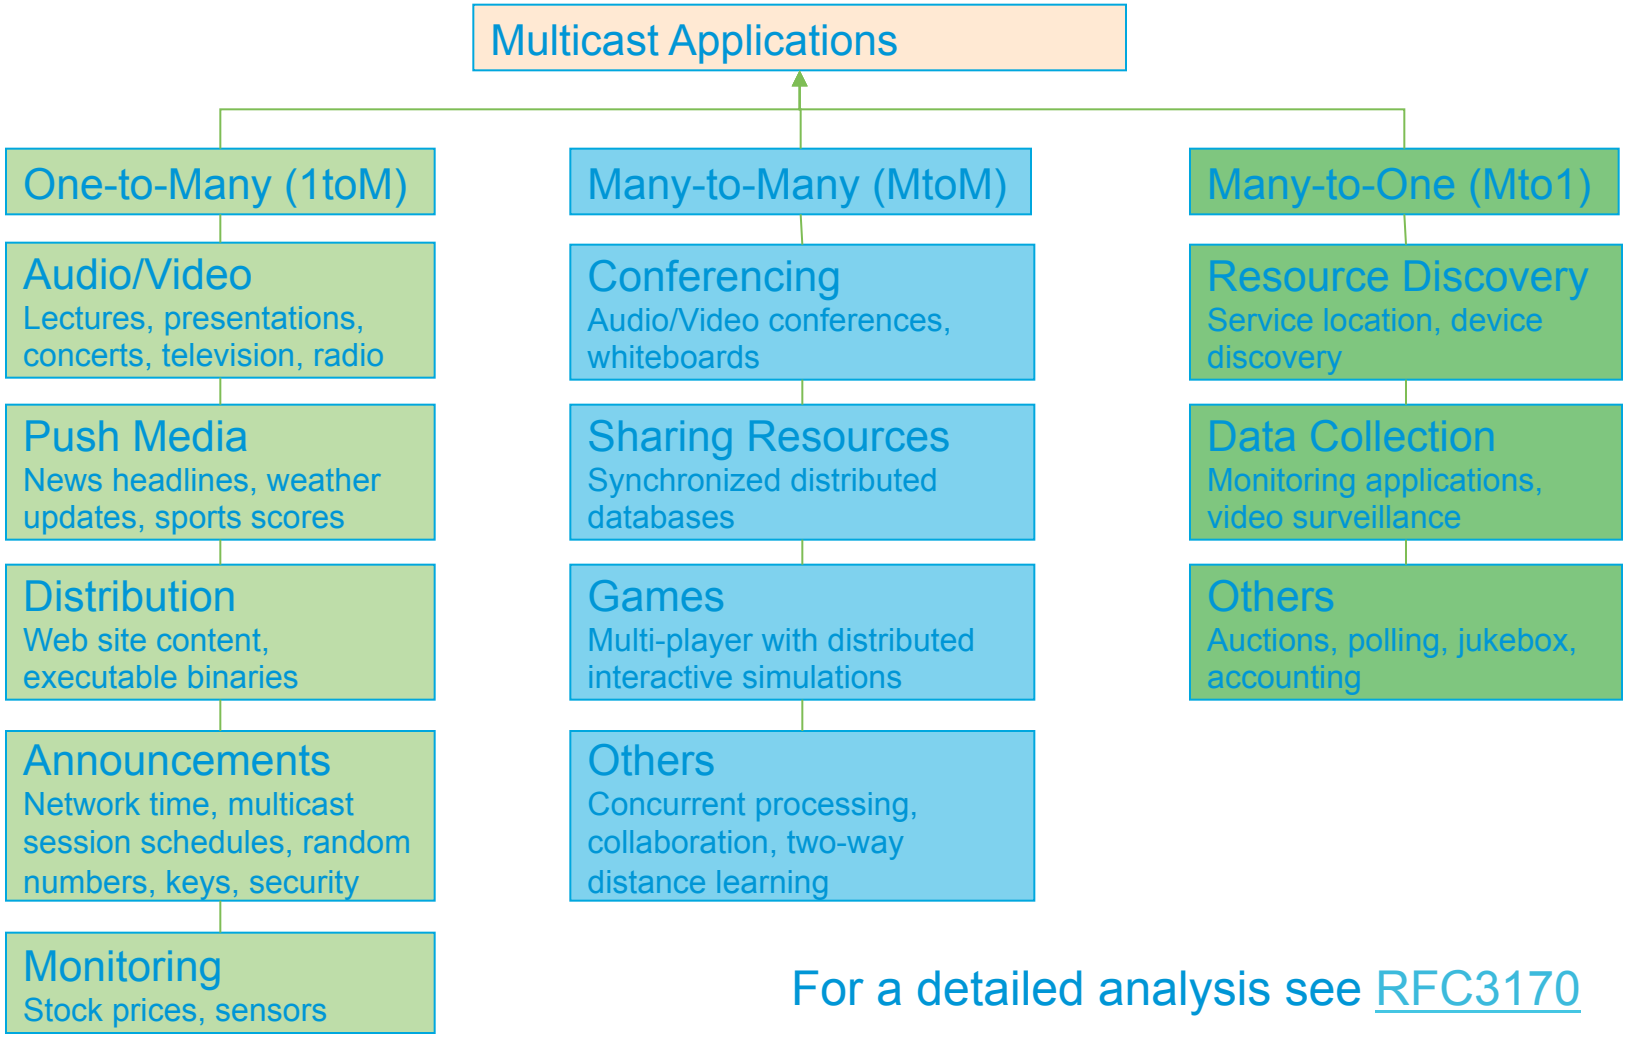

For a detailed analysis see RFC3170

Commercial Applications

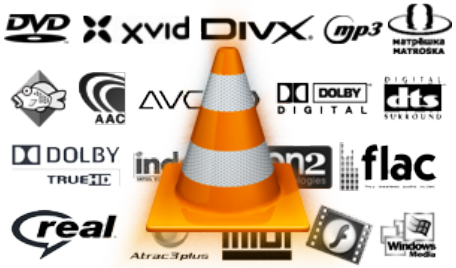

Microsoft Research
ConferenceXP
advanced collaboration • interactive distance learning

Multicast everywhere: Contribution, Distribution & Home Networks

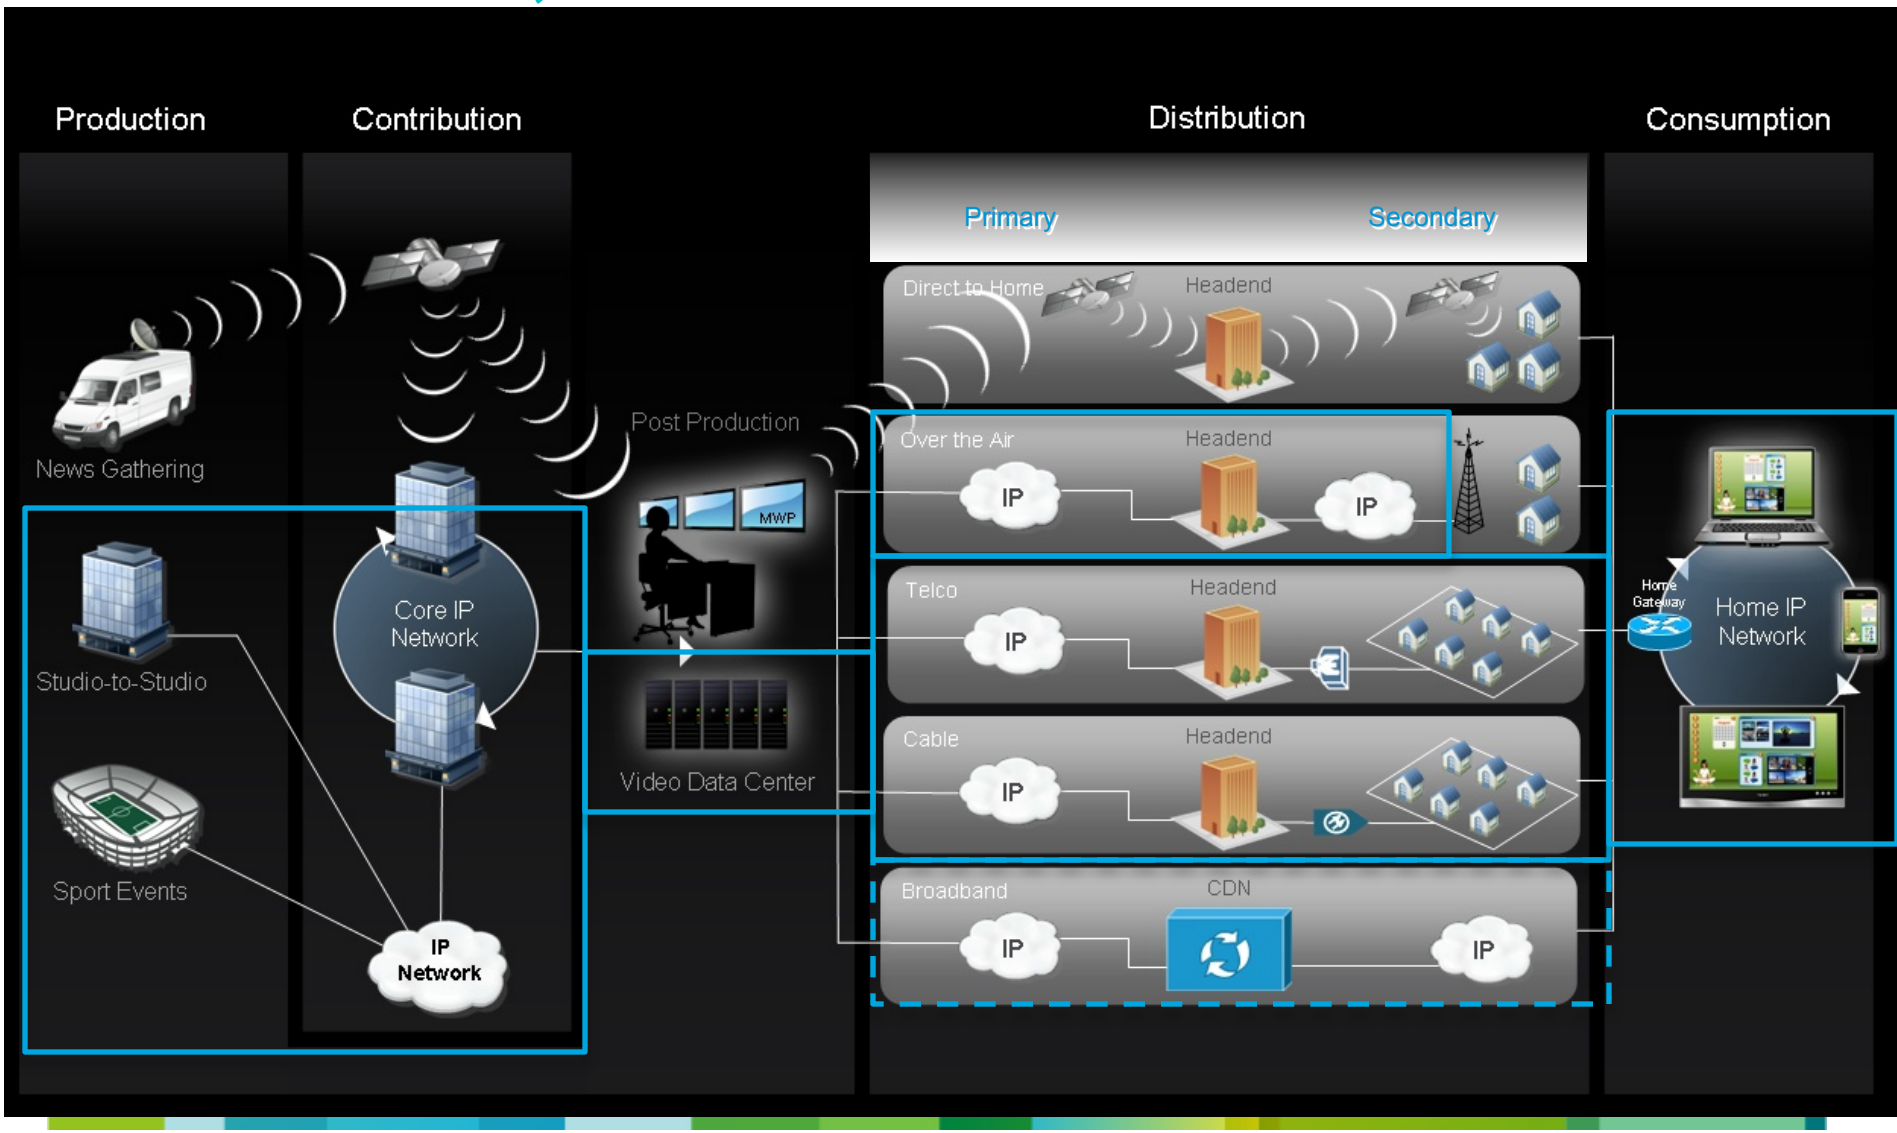

Multicast Group Addresses (224.0.0.0/4)

- Link-Local Address Range
  - 224.0.0.0–224.0.0.255
- Global Address Range
  - 224.0.1.0–238.255.255.255
  - Generally intended for “global” Internet scope multicast
- Administratively Scoped Address Range
  - 239.0.0.0–239.255.255.25
- Scope Relative Address Range
  - Top 256 addresses of a Scoped Address Range

Global Multicast Address Assignment - GLOP

- Static Group Address Assignment
  - RFC 3180—GLOP Addressing in 233.0.0.0/8
    - Group range: 233.0.0.0–233.255.255.255
      - Your Autonomous System number is inserted in middle two octets
      - Remaining low-order octet used for group assignment
    - EGLOP Addresses
      - Make use of private AS numbers
      - Assigned by a Registration Authority

Global Multicast Address Assignment - SSM

- Static Group Address Assignment
  - Source Specific Multicast
    - Address range: 232.0.0.0/8
    - Flows based on both Group **and** Source address
    - Two different content flows can share the same Group address without interfering with each other
      Better control and management of (re)sources
    - **Provides virtually unlimited address space!**
      - Preferred method for global one-to-many multicast

Private Administratively Scoped Address Assignment

- Assigned from the private 239.0.0.0/8 range
  - May be subdivided into geographic scopes ranges
  - Administration responsibility can be by scope range

Multicast Distribution Trees
Shortest Path/Source Tree

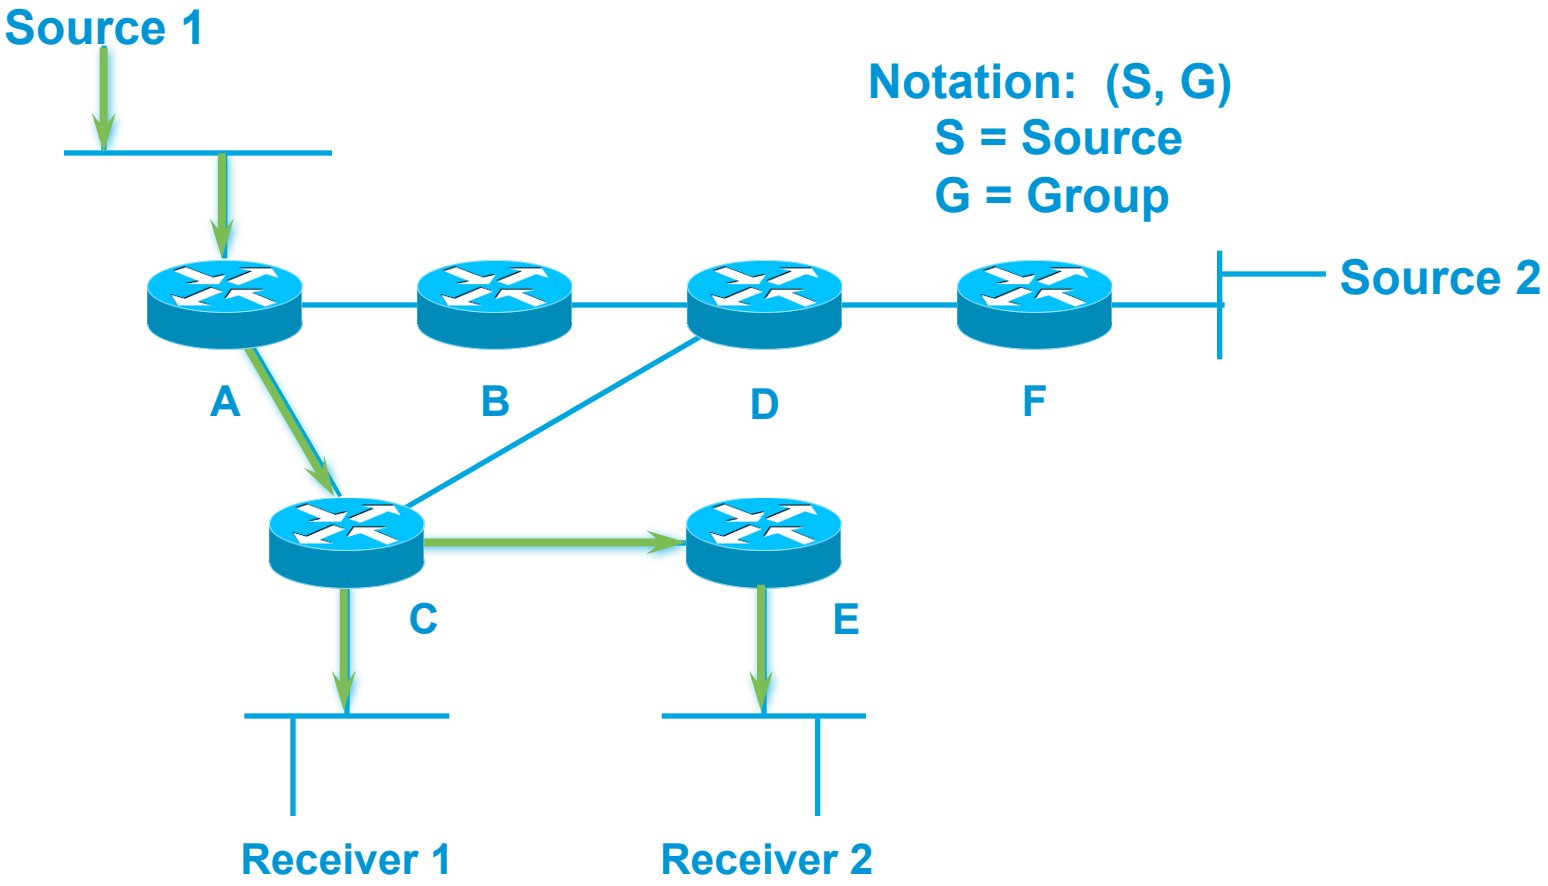

Multicast Distribution Trees
Shared tree

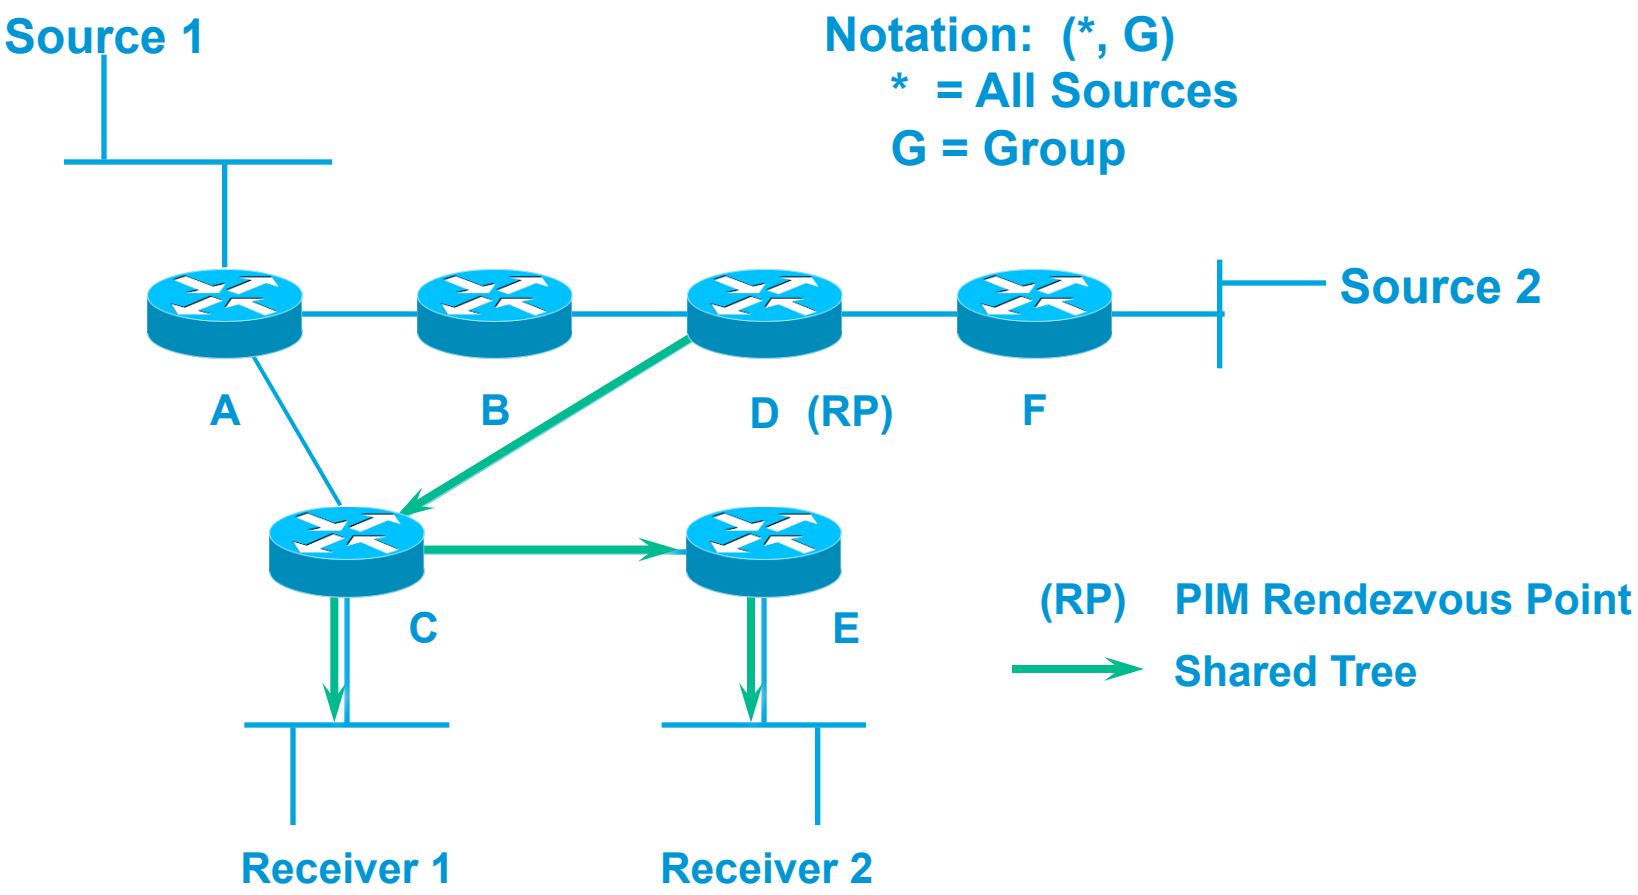

Multicast Distribution Trees
Shared tree and Source tree

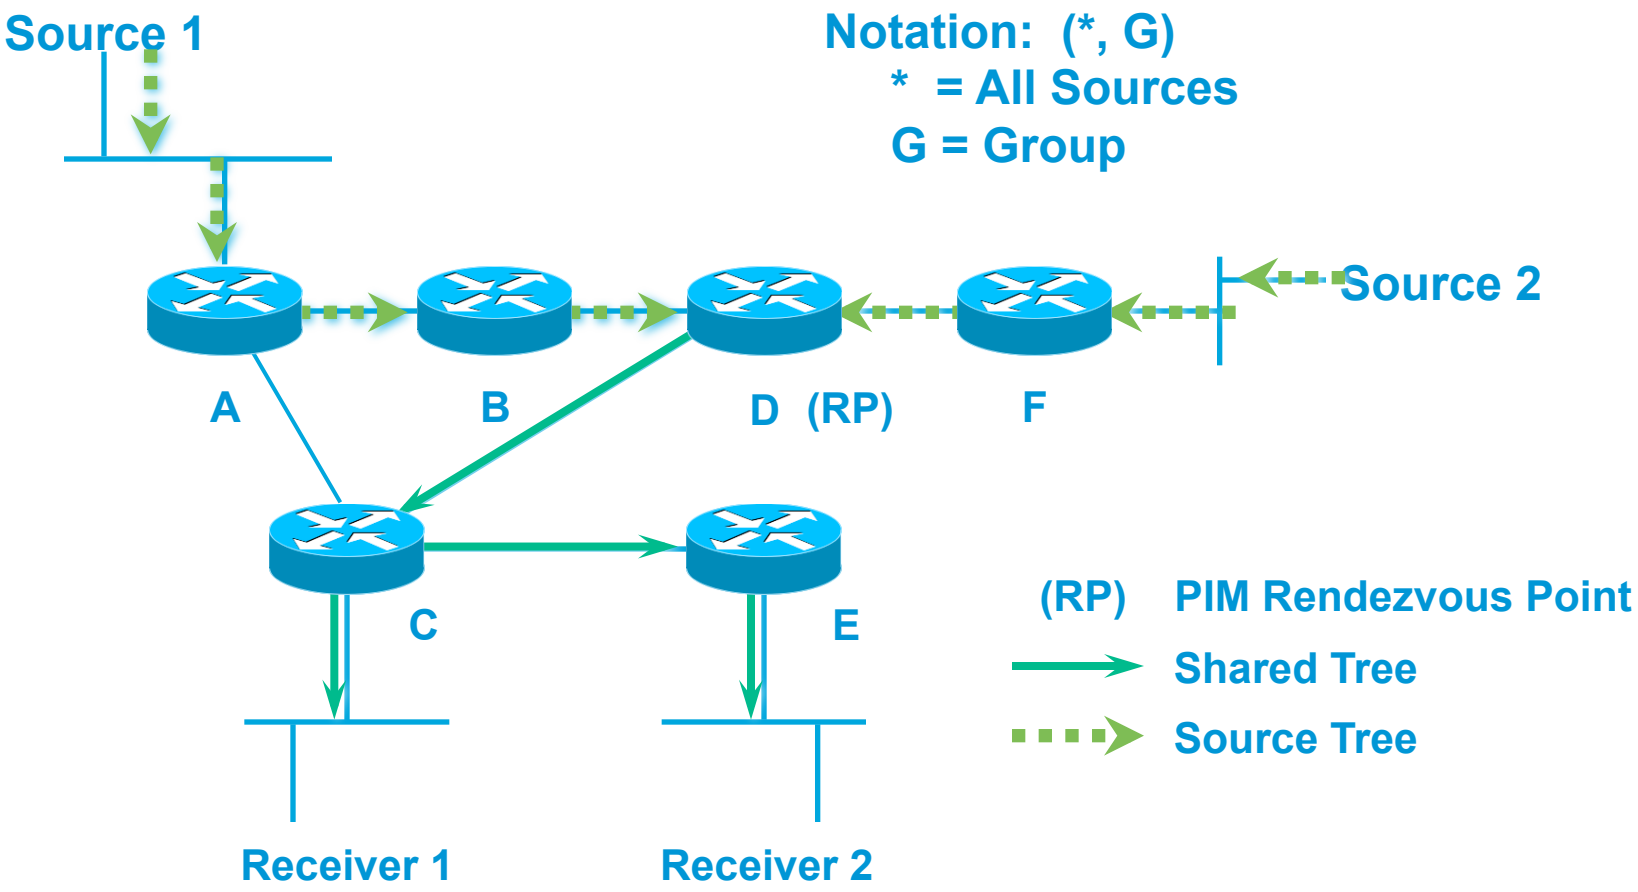

Barriers to Multicast Deployment

- Global Multicast Address Allocation
  - Dynamic Address Allocation
    - No adequate dynamic address allocation methods exist!
  - Static Address Allocation (GLOP)
    - Based on Autonomous System number
    - Insufficient address space for large Content Providers
- Multicast Content “Jammers”
  - Undesirable sources on a multicast group
    - “Capt. Midnight” sources bogus data/noise to group
    - Can cause Denial of Service attack by congesting low speed links

Source Specific Multicast (SSM)

- Uses Source Trees only
- Assumes one-to-many model
  - Most Internet multicast fits this model
  - Video Contribution and Distribution also fit this model
- Hosts responsible for source discovery
  - Typically via some out-of-band mechanism
    - Web page, Content Server, Management tool, etc.
  - Eliminates need for Rendezvous Points and Shared Trees
  - Eliminates need for MSDP (inter domain RP based multicast)

SSM Overview

- Hosts join a specific source within a group
  - Content identified by specific (S,G) instead of (*,G)
  - Hosts responsible for learning (S,G) information
- Last-hop router sends (S,G) join toward source
  - Shared Tree is never Joined or used
  - Eliminates possibility of content Jammers
  - Only specified (S,G) flow is delivered to host
- Eliminates Networked-Based Source Discovery
  - No RPs for SSM groups
- Simplifies address allocation
  - Content sources can use same group without fear of interfering with each other

SSM Example

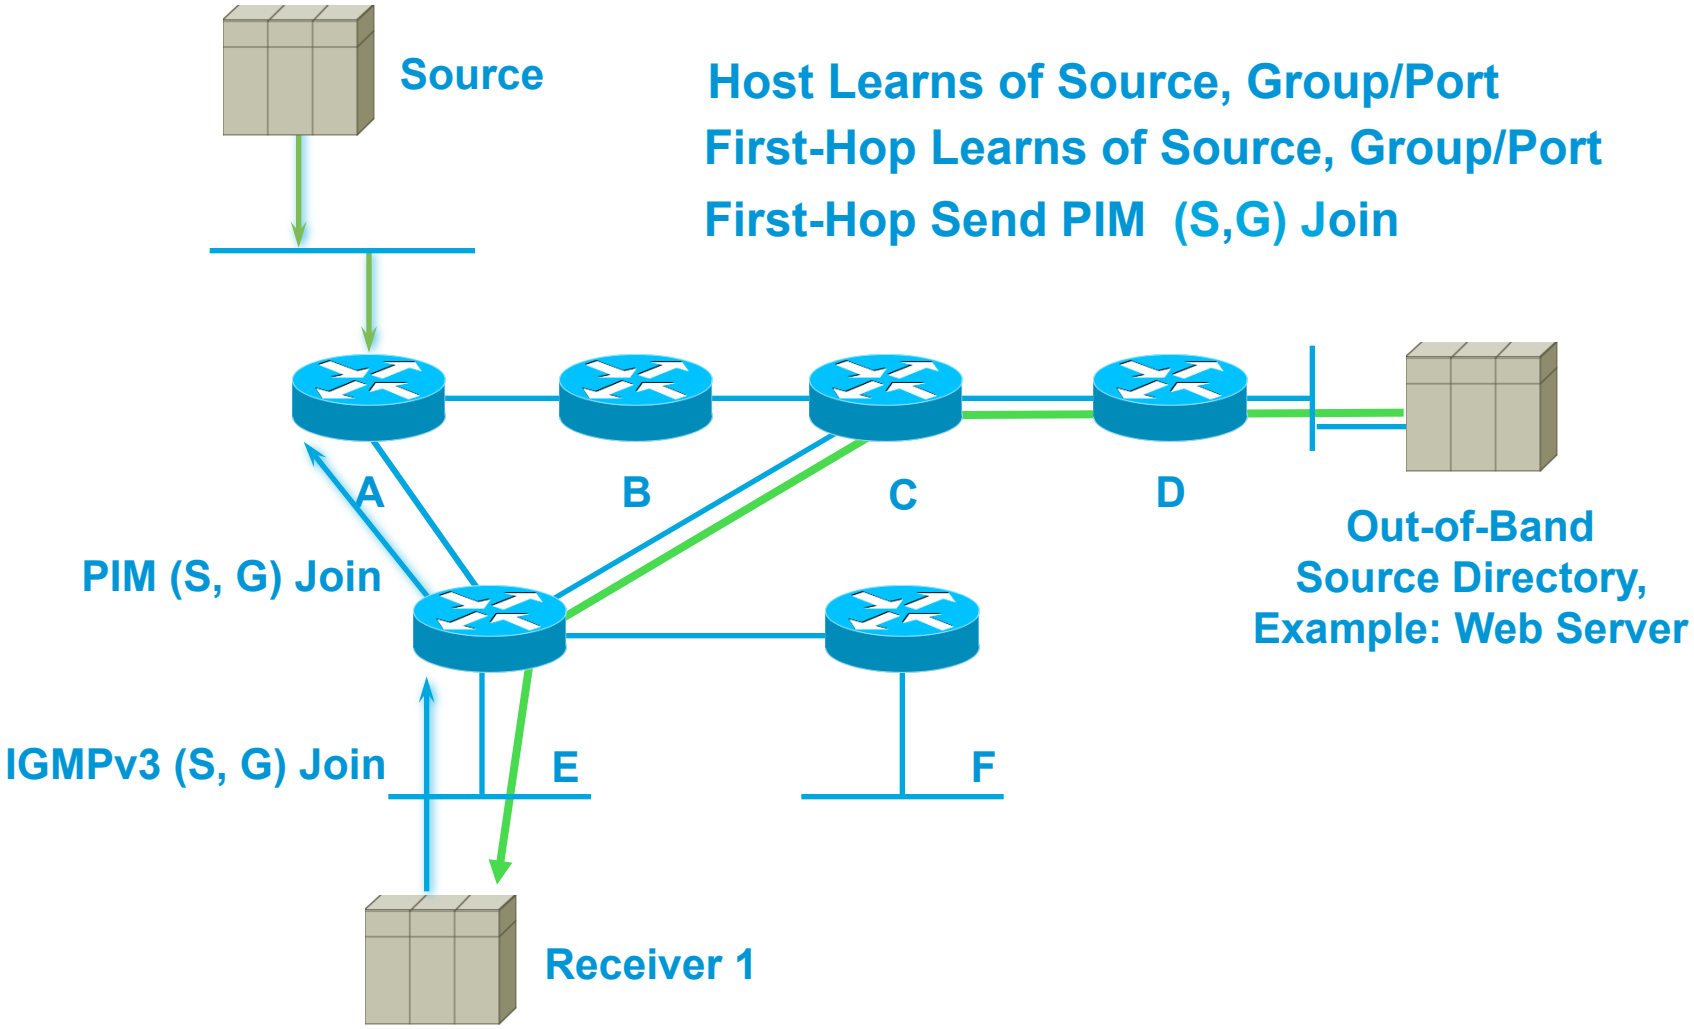

SSM Example

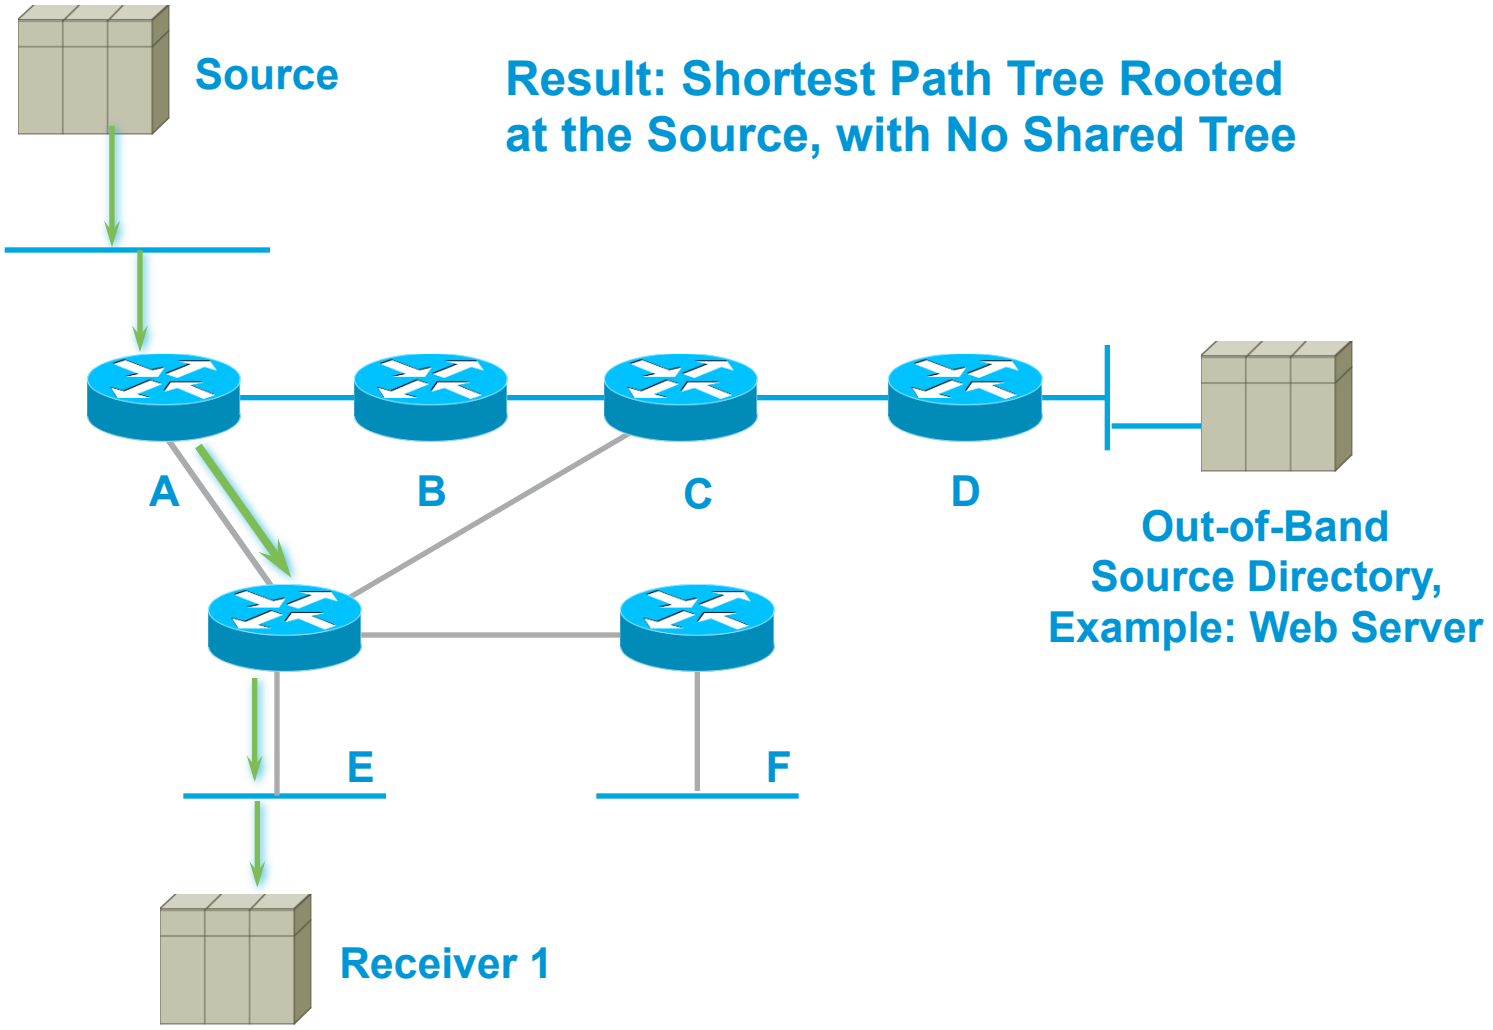

IGMP

- How hosts (receivers) tell routers about group membership
- Routers solicit group membership from directly connected hosts
- RFC 1112 specifies first version of IGMP
- RFC 2236 specifies IGMPv2
  - Most widely deployed and supported
- RFC 3376 specifies IGMPv3
  - Growing support (required for SSM)

IGMPv3 Example

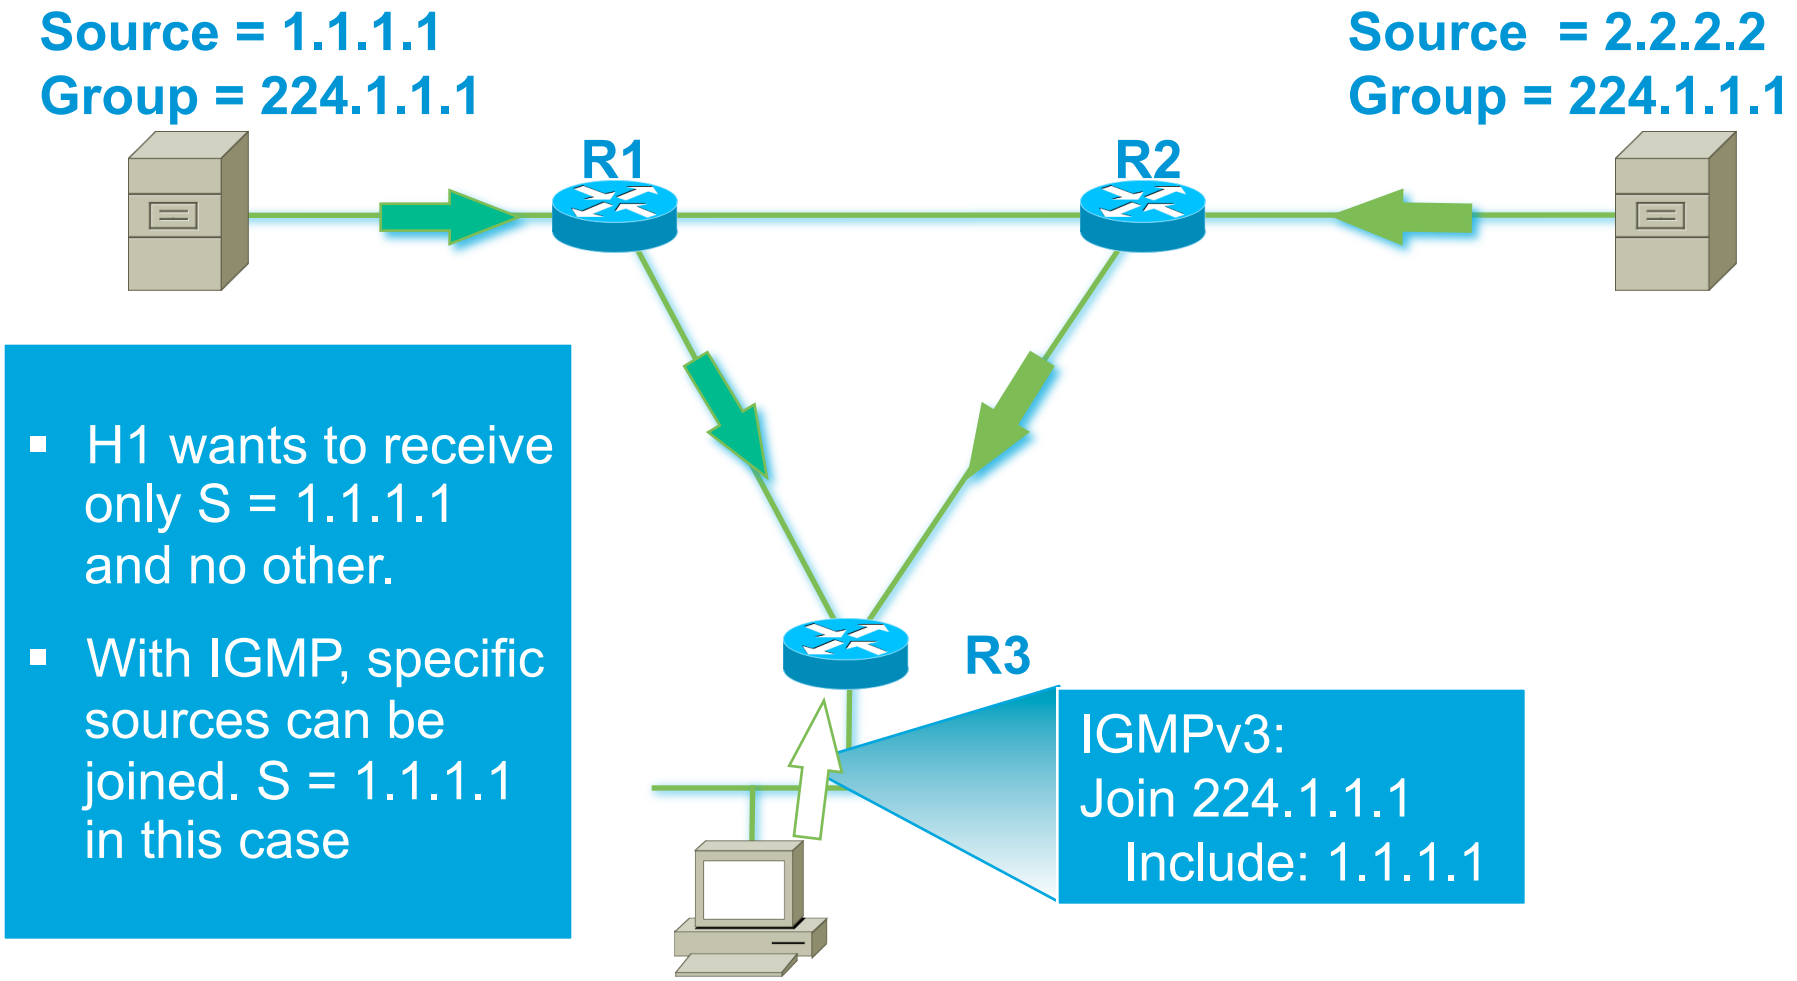

- H1 wants to receive only S = 1.1.1.1 and no other.
- With IGMP, specific sources can be joined. S = 1.1.1.1 in this case

IPv6 and Multicast

- Designed into IPv6 specification from the start
  - Multicast was originally designed after the initial IPv4 protocol
- Natively part of discovery and information exchange between protocols
  - Link local discovery
  - Routing protocols message exchange between neighbors
  - ICMPv6 messages
  - Service discovery (DNS, NTP, …)
- Multicast in IPv6 is part of the system and must be treated as such
- Leverage the rollout of IPv6 to ensure multicast is part of the service!

IPv6 Multicast Addresses - RFC 4291

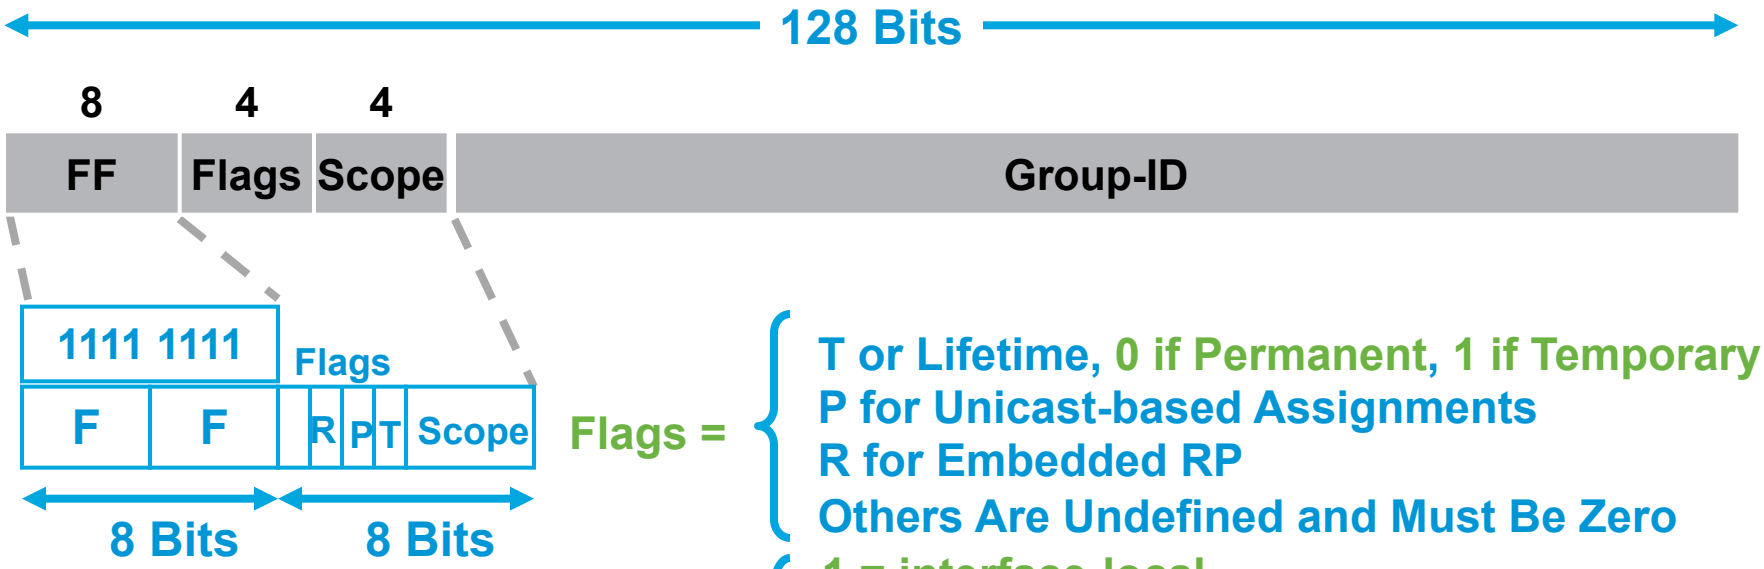

Note: other scopes (6, 7, 9-D) are unassigned but can be used

IPv6: Multicast Listener Discovery – MLD
Multicast Host Membership Control

- MLD specified in RFC's 2710 and 3810
- MLD is equivalent to IGMP in IPv4
- MLD messages transported over ICMPv6
- MLD uses link local source addresses
- MLD packets use "Router Alert" option in IPv6 Hop-by-Hop extension header (RFC 2711) with Hop Limit =1

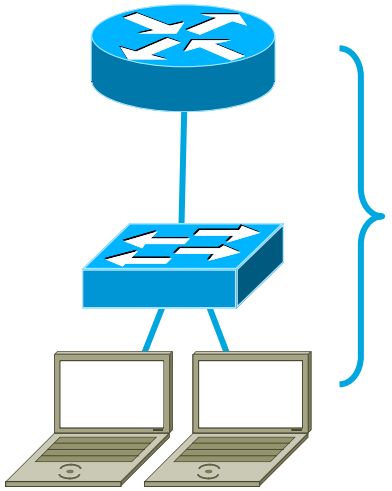

IPv6: Multicast Listener Discovery – MLD
Multicast Host Membership Control

- Version number “confusion”:

  MLDv1 (RFC 2710) ~ IGMPv2

  MLDv2 (RFC 3810) ~ IGMPv3

- MLDv2 router compatible with MLDv1 hosts
- SSM transition through SSM mapping for MLDv1 messages – static or DNS
- MLD snooping

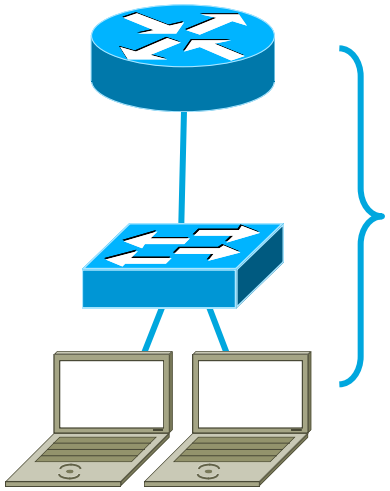

IPv6 Multicast Based Multimedia Services (NTT-East)

- NTT-East rolled out native IPv6 multicast services instead of IPv4 offering IPTV, music and games:

  http://www.ipv6style.jp/en/action/20040902/index.shtml

  http://www.networkworld.com/news/2009/010809-ntt-ipv6-tv.html

Growth of Internet based video services

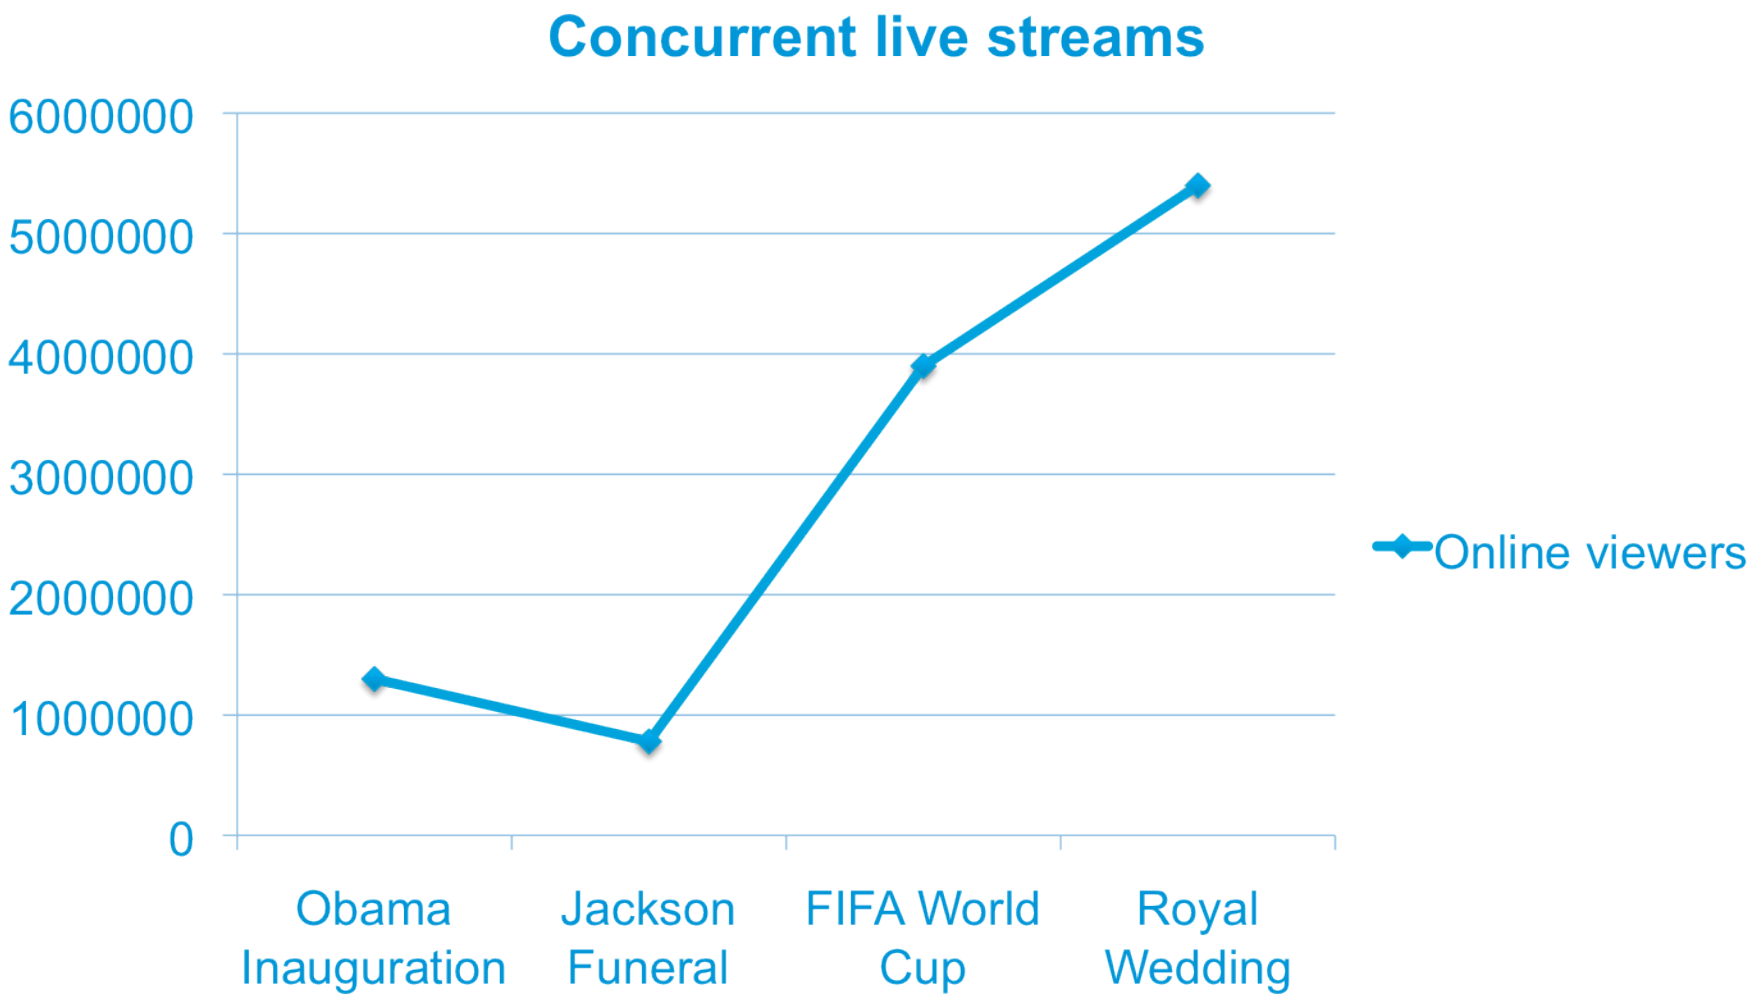

Growth of Internet based video services

- Live video:
  - 4% of all consumer traffic by end 2010
  - 9% of all video traffic in 2014
- 2009:
  - 107 PB/month
  - 2775 PB/month
- 2014:
  - 4075 PB/month in 2014
  - 19468 PB/month in 2014

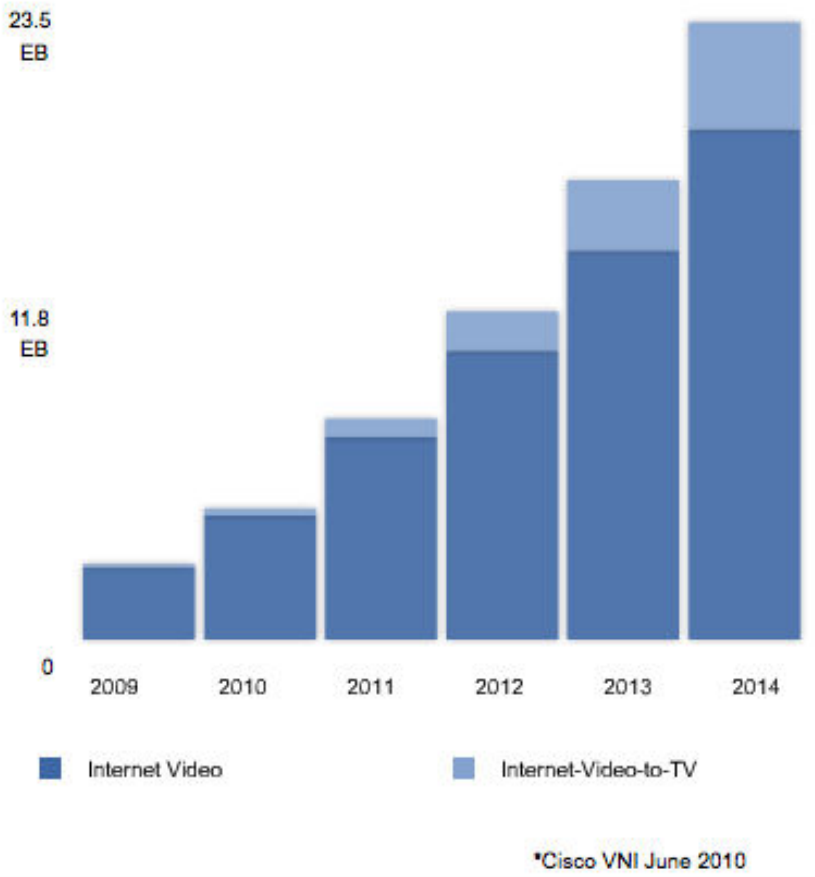

Source: Cisco Visual Networking Index

Multicast attempt across Unicast "islands"

Even though Content Owner, Core Provider are Multicast-Enabled, many Edge Networks are Unicast-Only

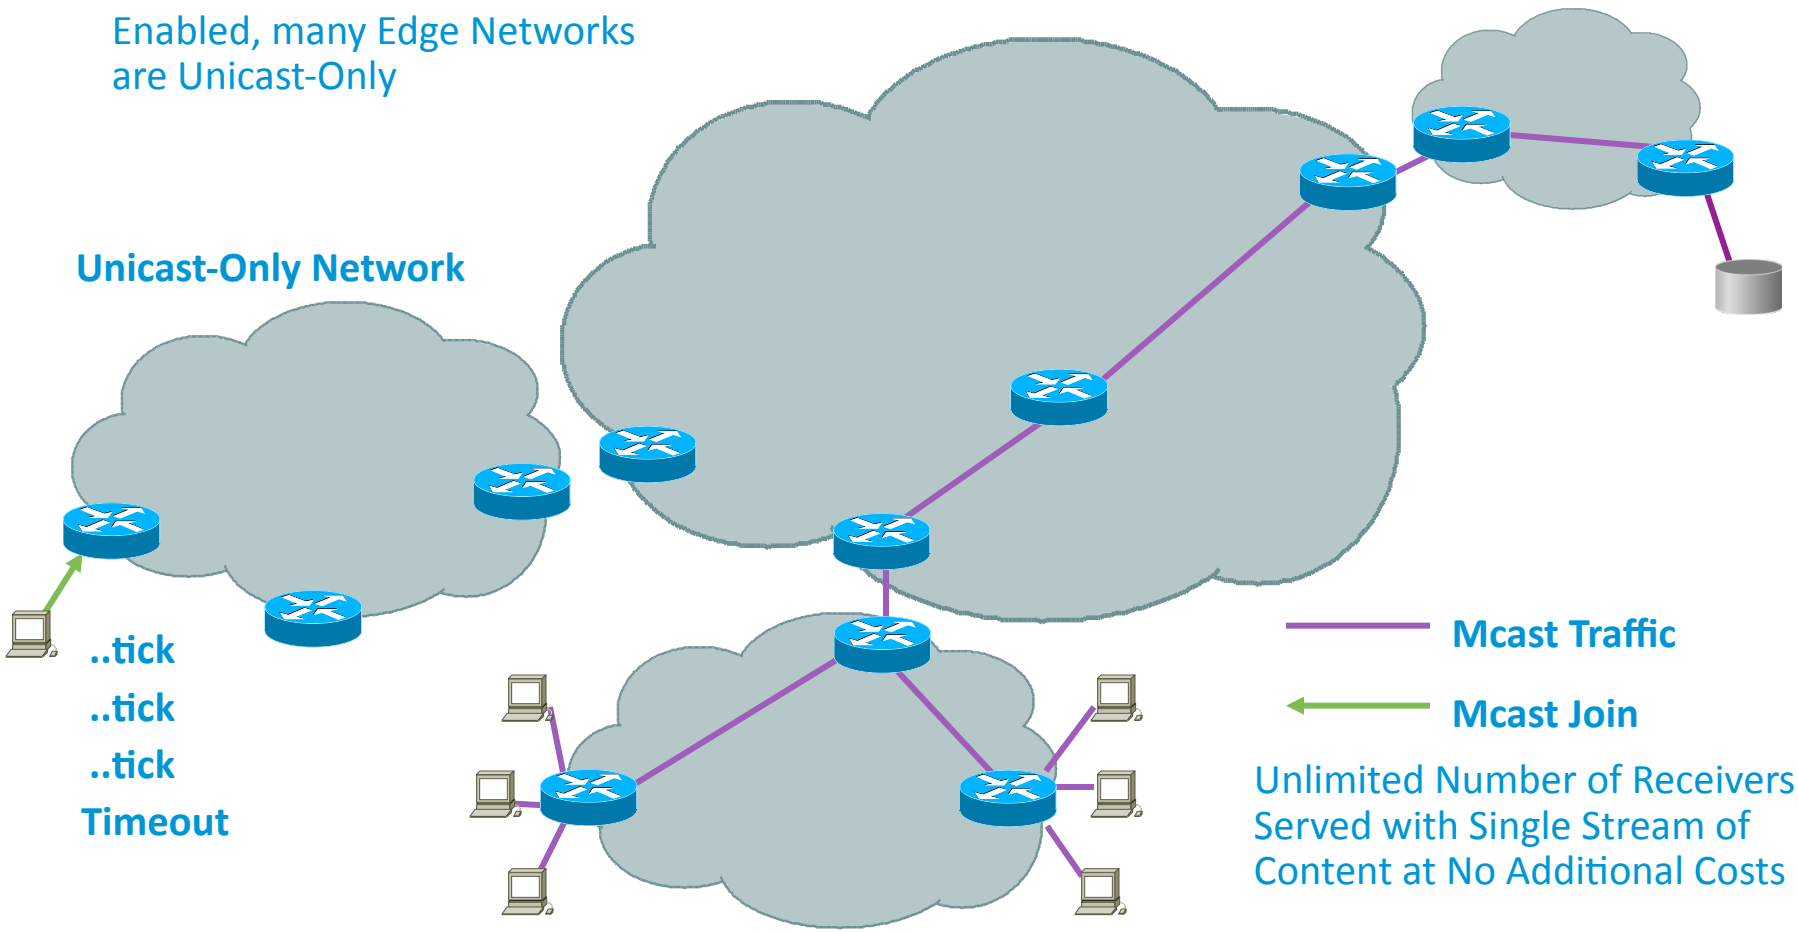

Unicast fallback

More Receivers consume more Resources and costs Content owner $$$

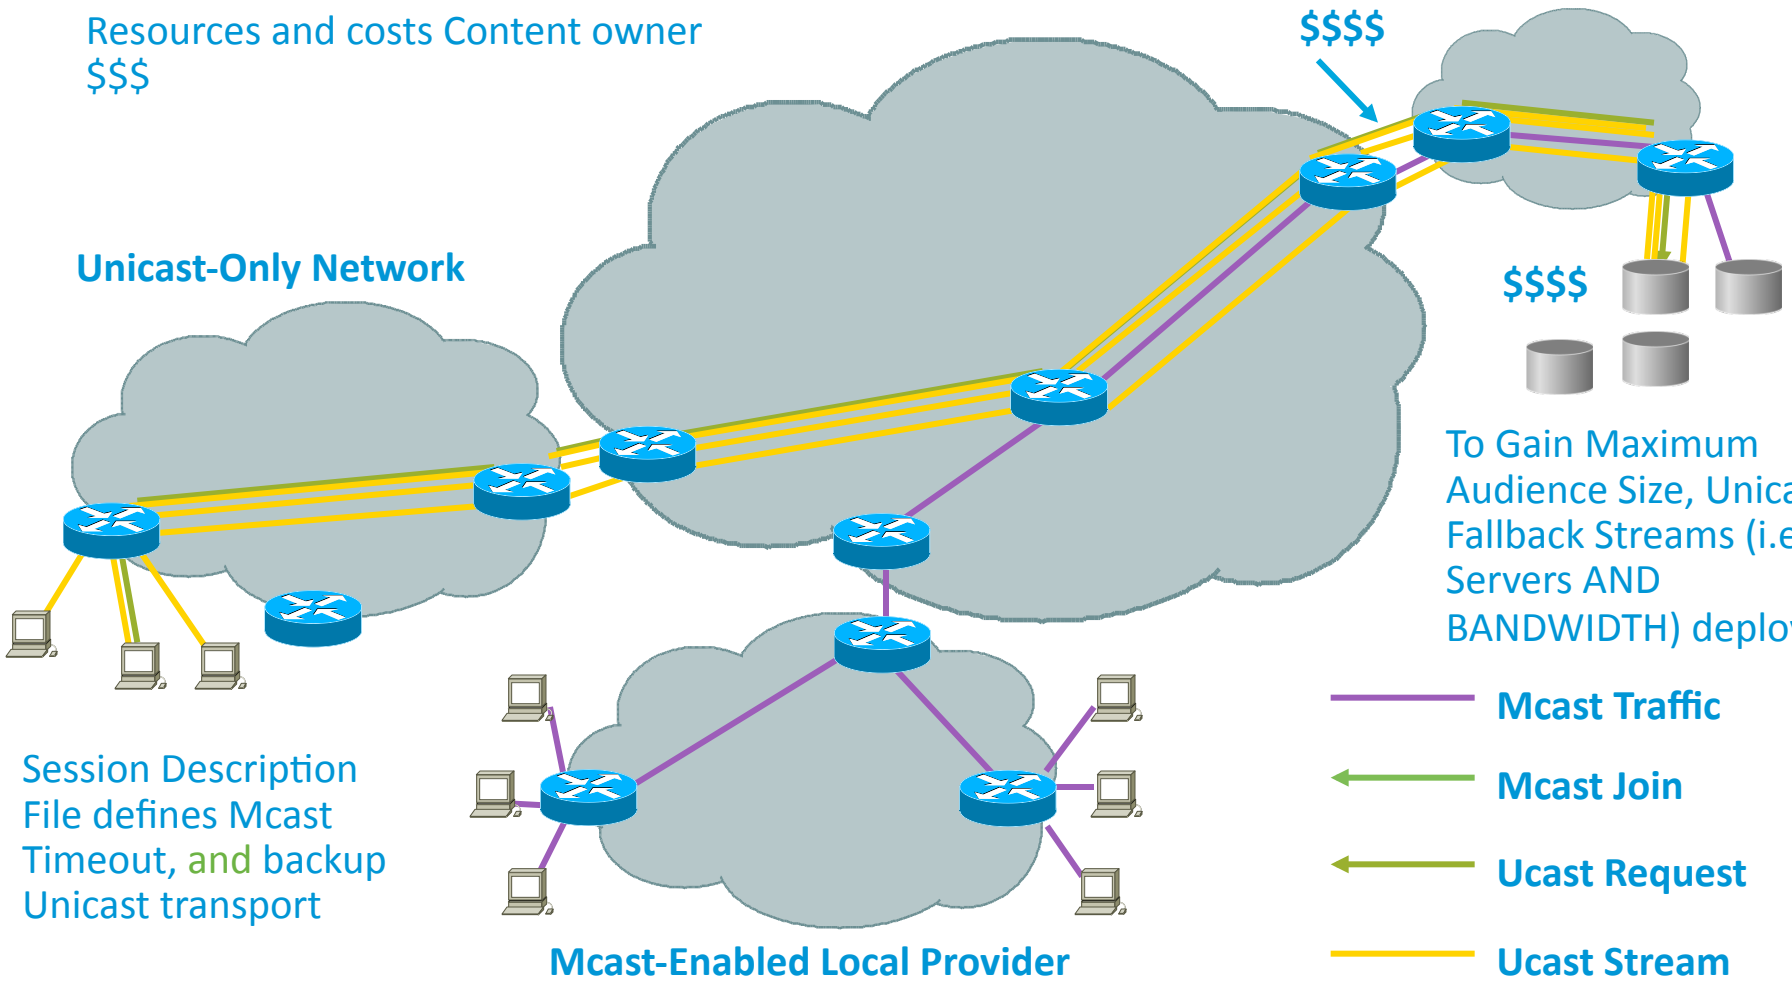

What's Wrong?

- Multicast in the Internet is an all or nothing solution
  - Each receiver must be on an IP multicast-enabled path
  - Many core networks have IP multicast-enabled, but few edge networks accept multicast transit traffic
- Even multicast-aware content owners are forced to provide unicast streams to gain audience size
- Unicast is unable to scale for streaming live content at current growth rates
  - Splitters/caches just distribute the problem
    - Still has a cost per user
  - As receiver bandwidth increases, problem gets worse

Automatic IP Multicast Tunneling

- Automatic IP Multicast Tunneling:

  http://tools.ietf.org/id/draft-ietf-mboned-auto-multicast

- Designed to provide a migration path to a fully multicast enabled backbone
- Allows multicast to reach unicast-only receivers without the need for any explicit tunneling
- Provide benefits of multicast wherever multicast is already deployed

  Hybrid solution

  Multicast networks get the benefit of multicast

- Works seamlessly with existing applications

  Requires only client-side shim (somewhere in client) and router support (in some places)

AMT architecture

- AMT Gateway
  - Initiates connection to the multicast network via an AMT Discovery message
  - Discovery message sent to “well known” Anycast address
  - May be a host (PC, Mac, Xbox, Android, …) or a gateway/router
- AMT Relay
  - Listens for AMT Discovery messages to build AMT tunnel to requesting Gateways
  - May be on a router at the unicast/multicast boundary or in an appliance near the boundary

AMT in action – discovering the Relay

Anycast Address allows all Gateways to find "Closest" Relay— nearest edge of Source Multicast Topology

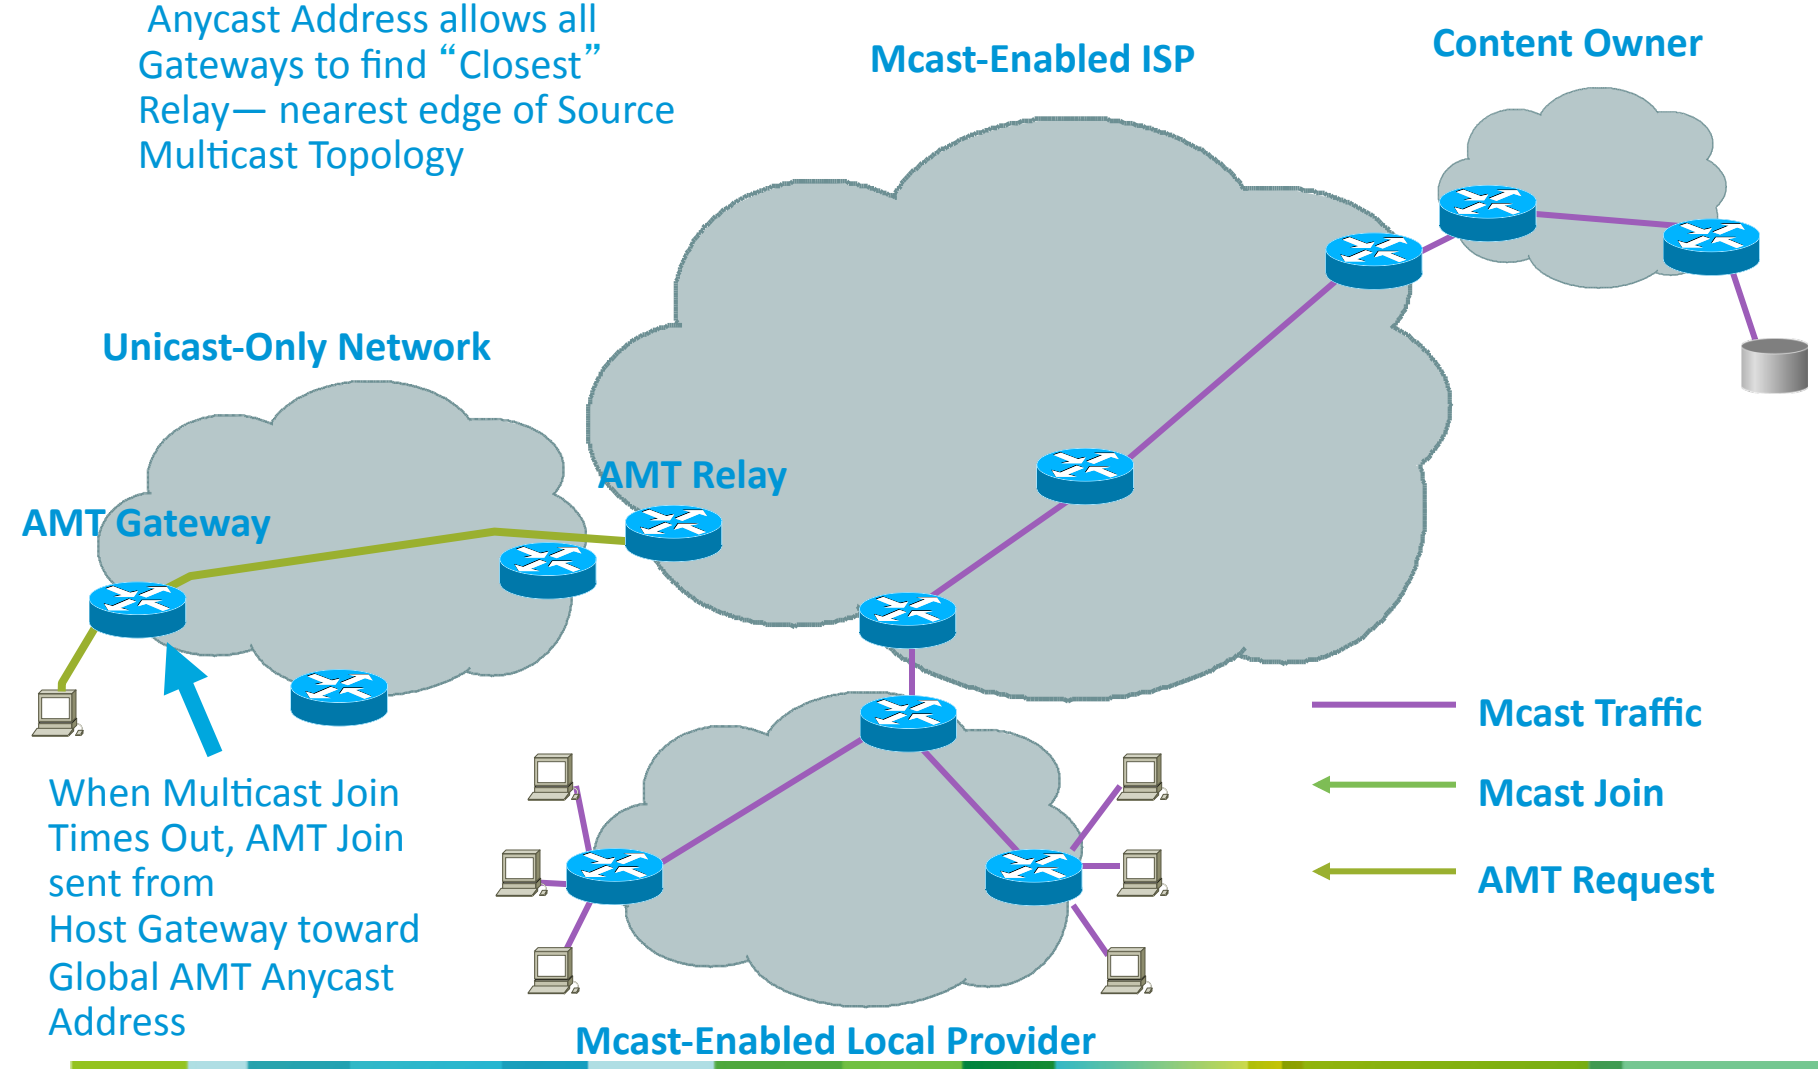

When Multicast Join Times Out, AMT Join sent from Host Gateway toward Global AMT Anycast Address

AMT in action – building the tunnel

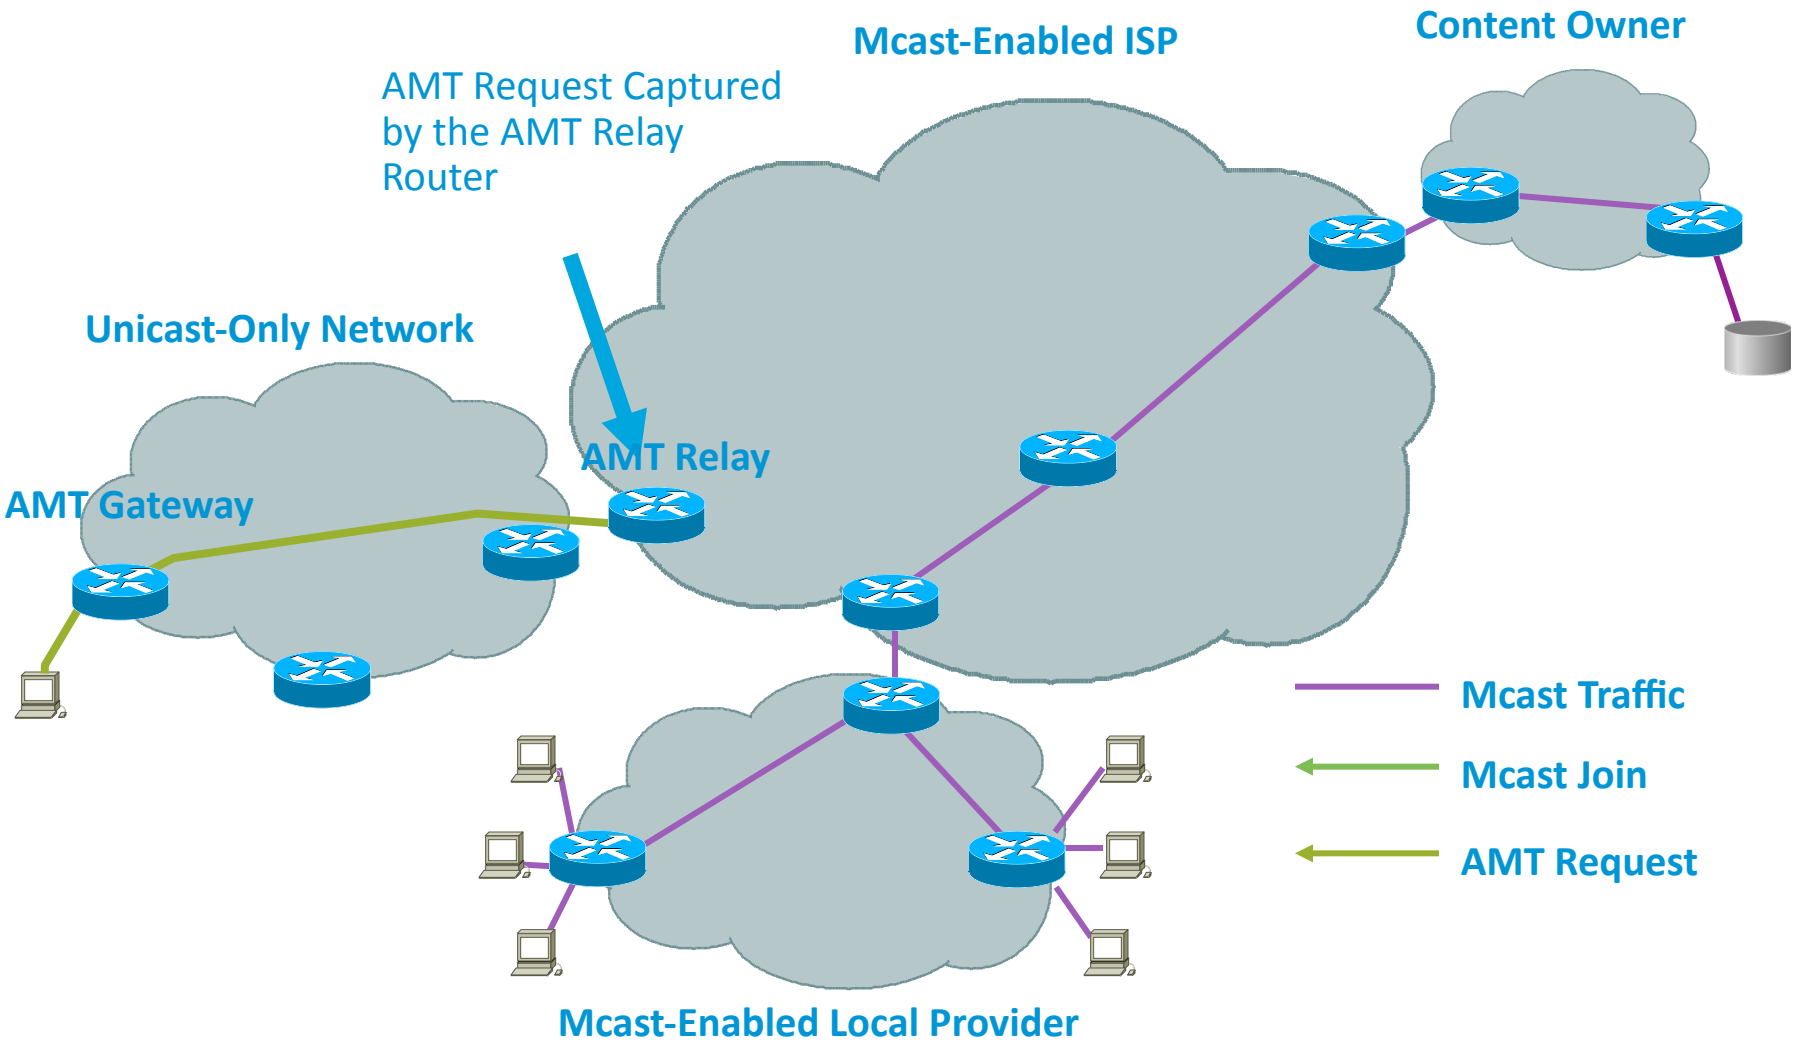

AMT in action – joining the source stream

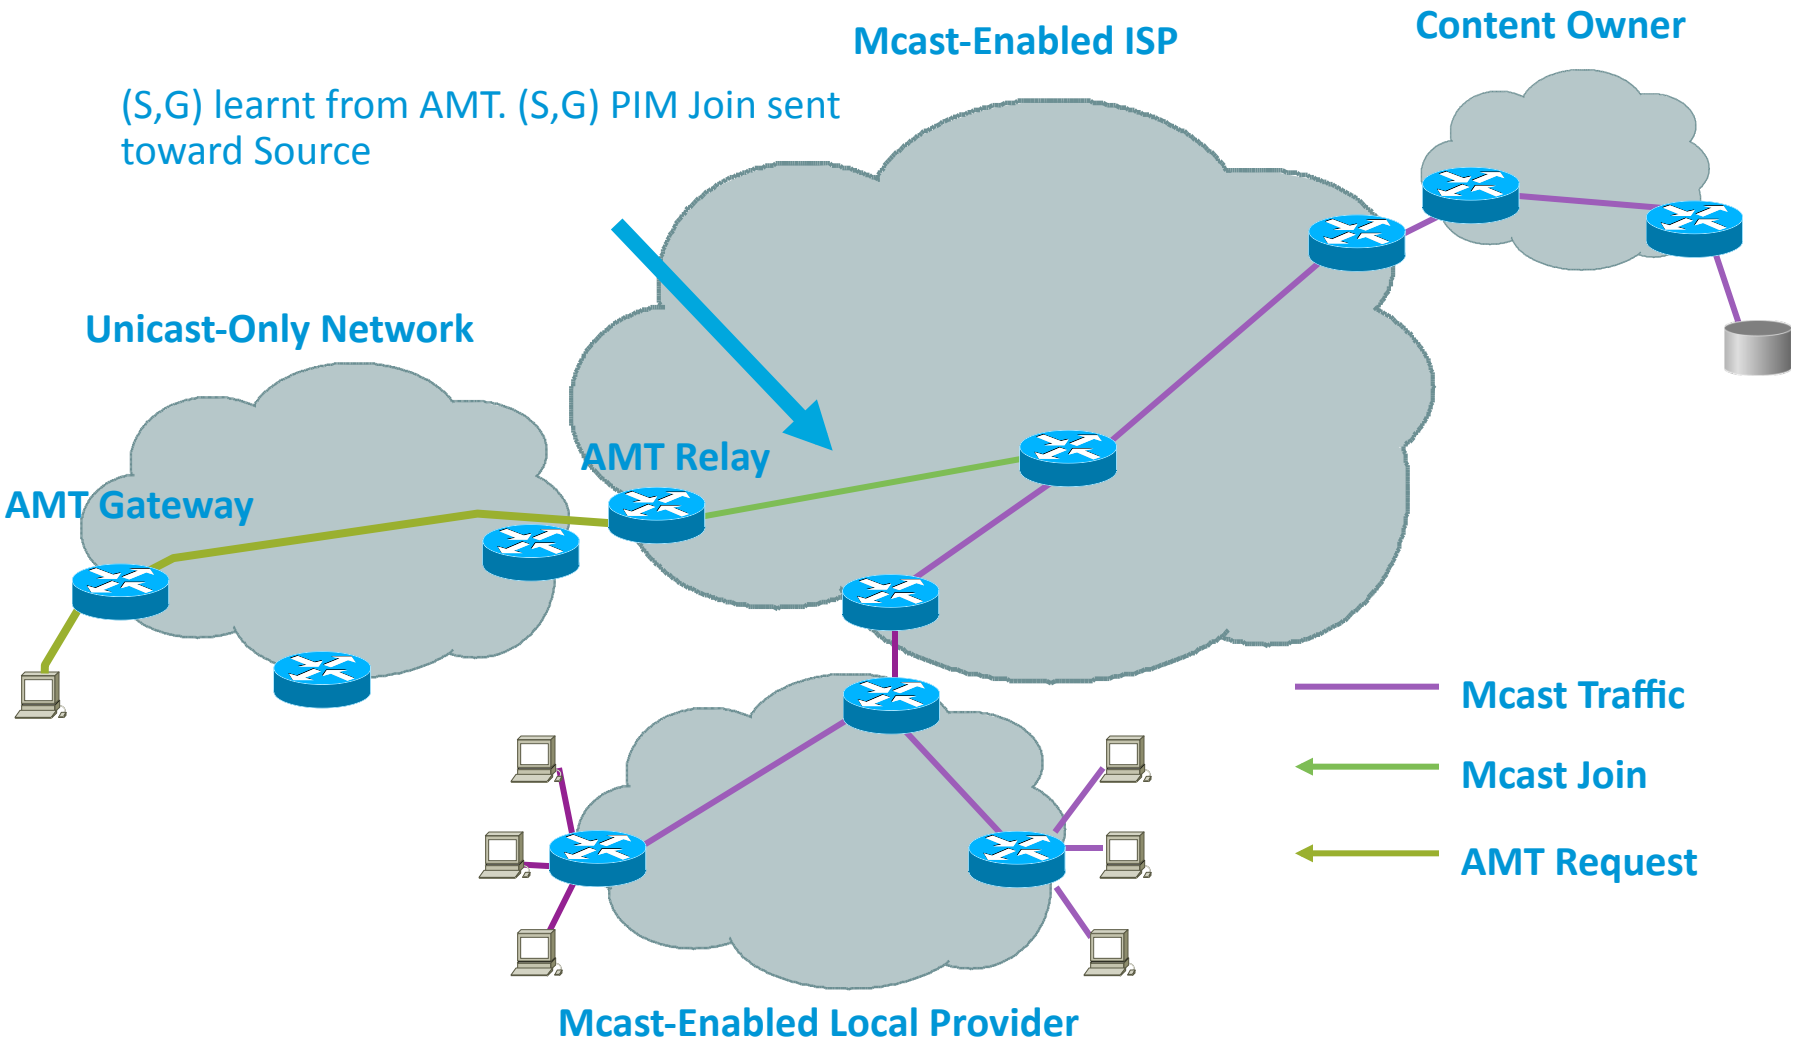

AMT in action – serving single receiver

Relay Replicates to downstream receivers, adding unicast header (destined to Receiver)

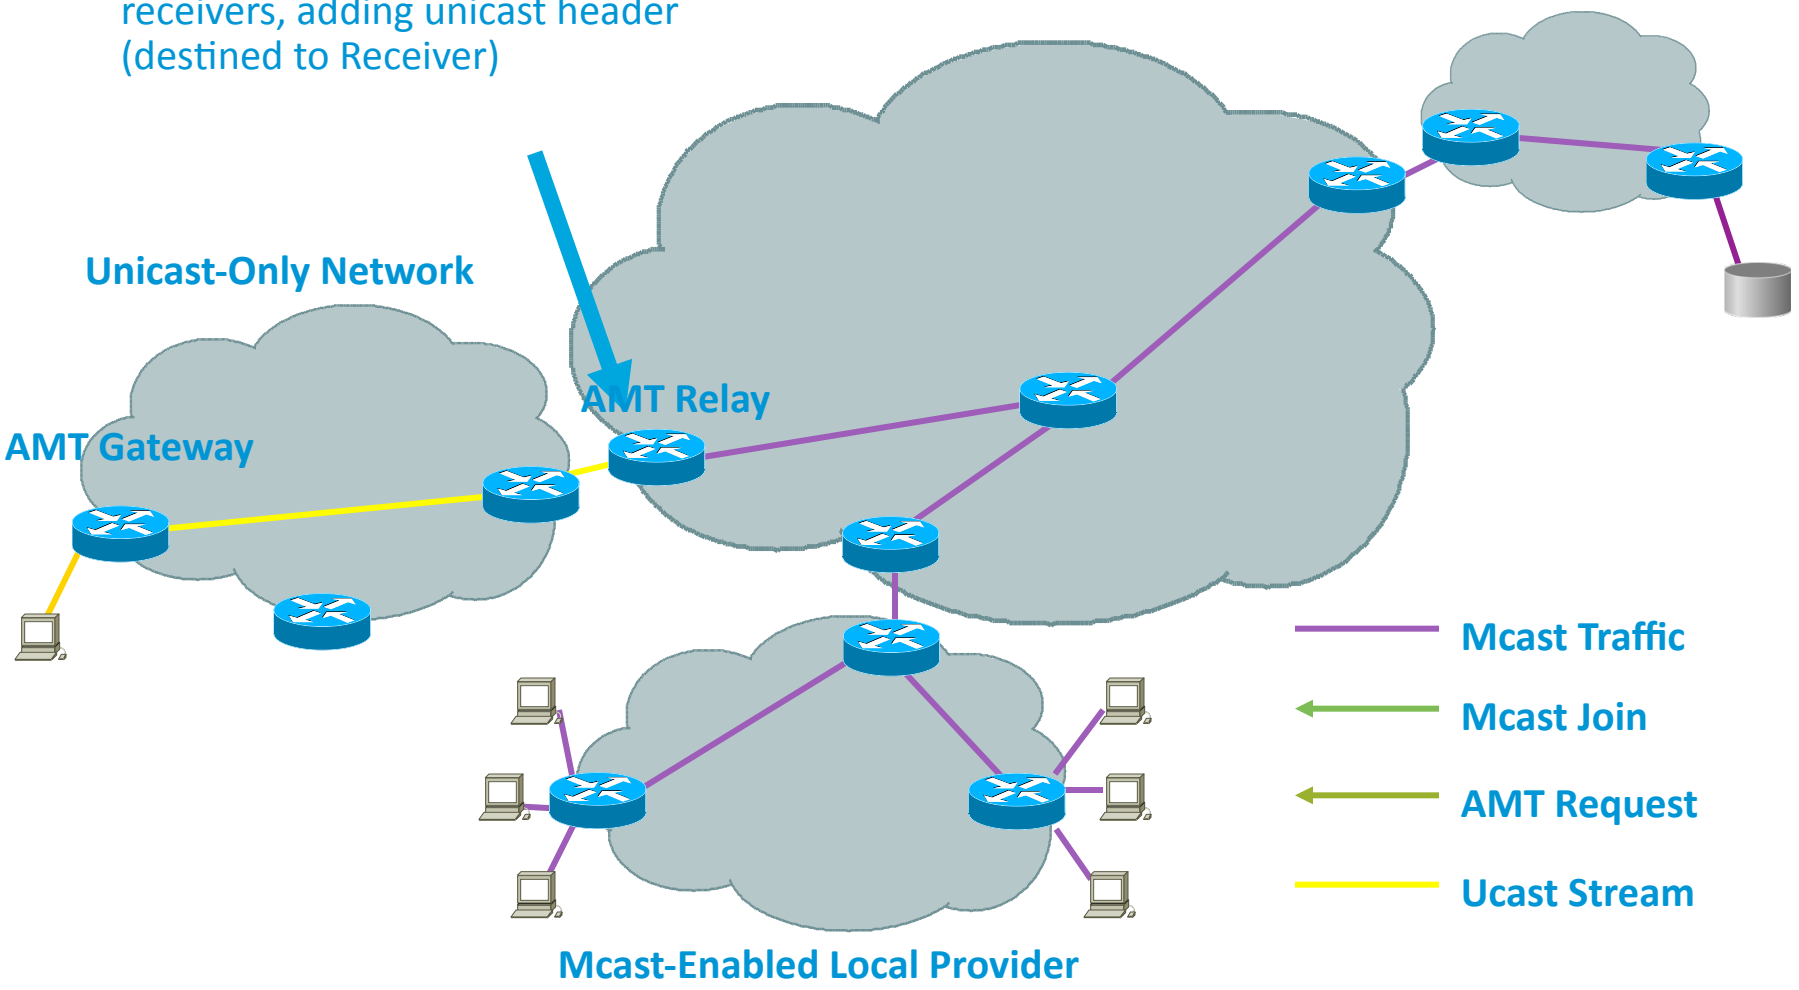

AMT in action – serving multiple receivers

Additional receivers served by single Relay. Benefits of Multicast retained by Content Owner and all enabled Networks in path

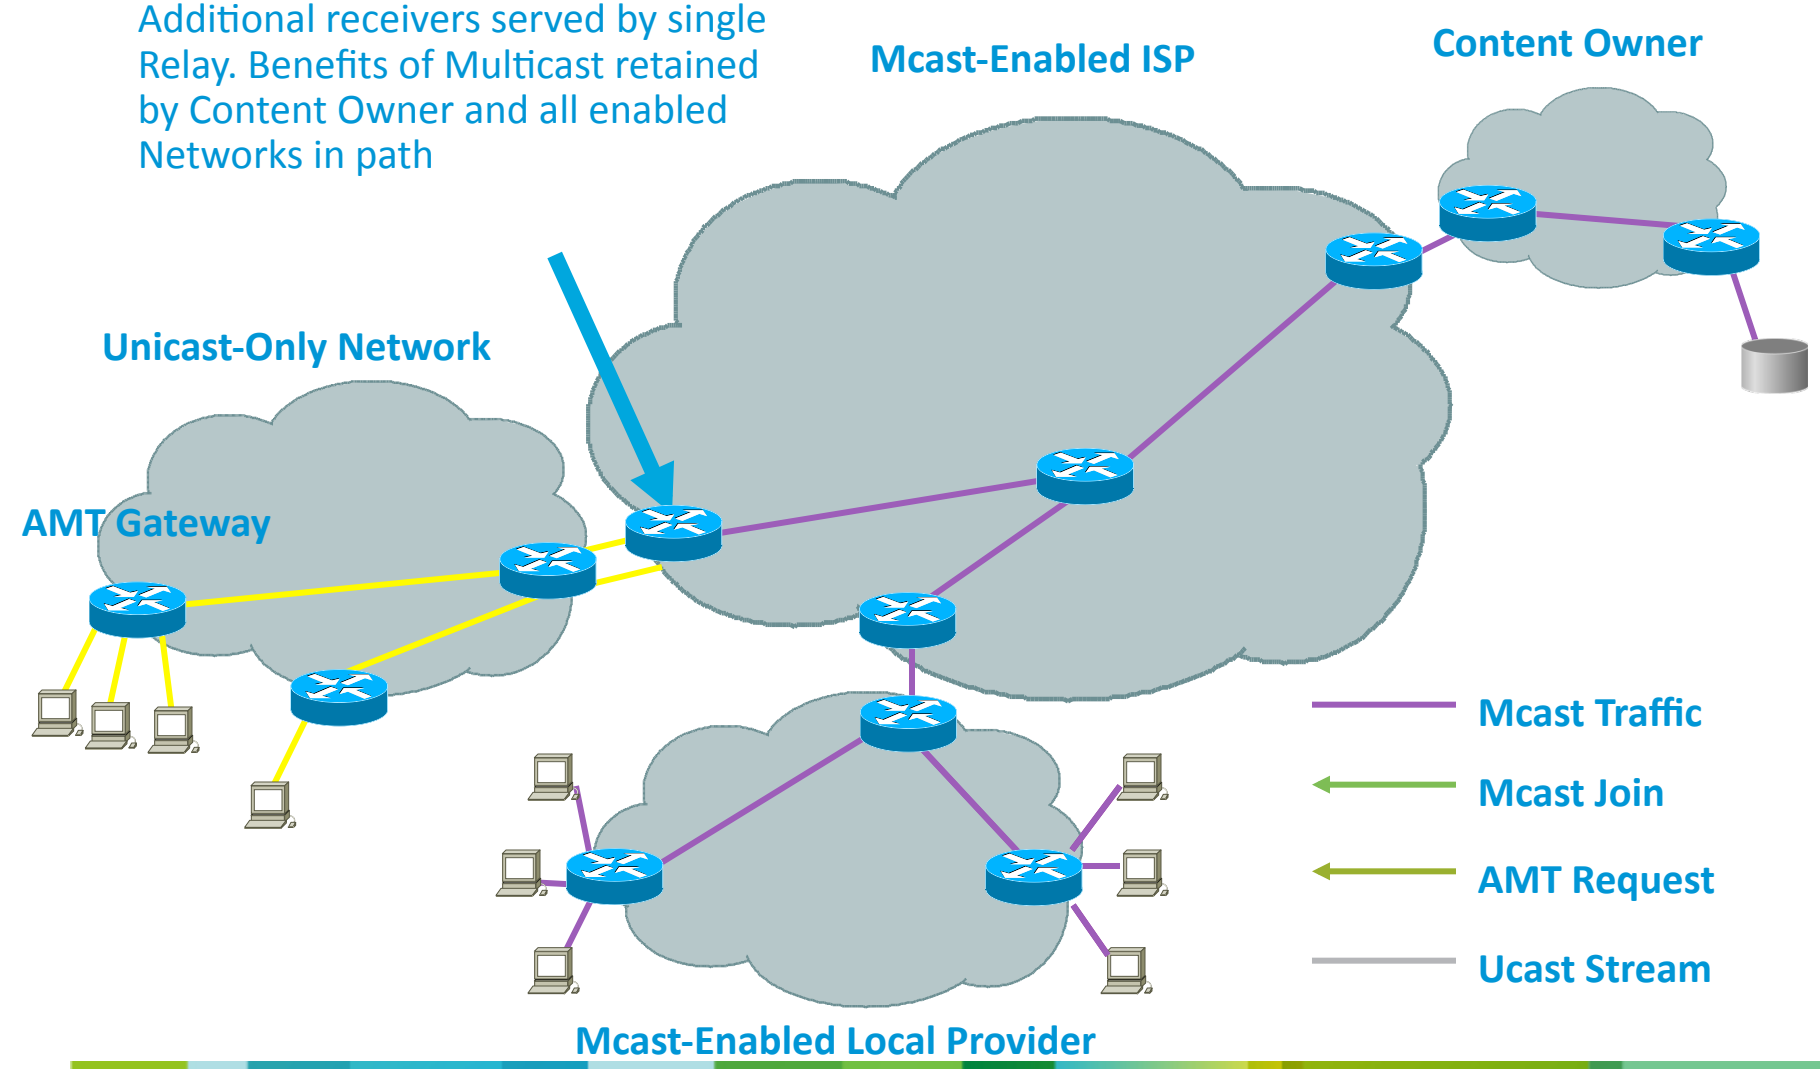

AMT in action – “spreading the word”

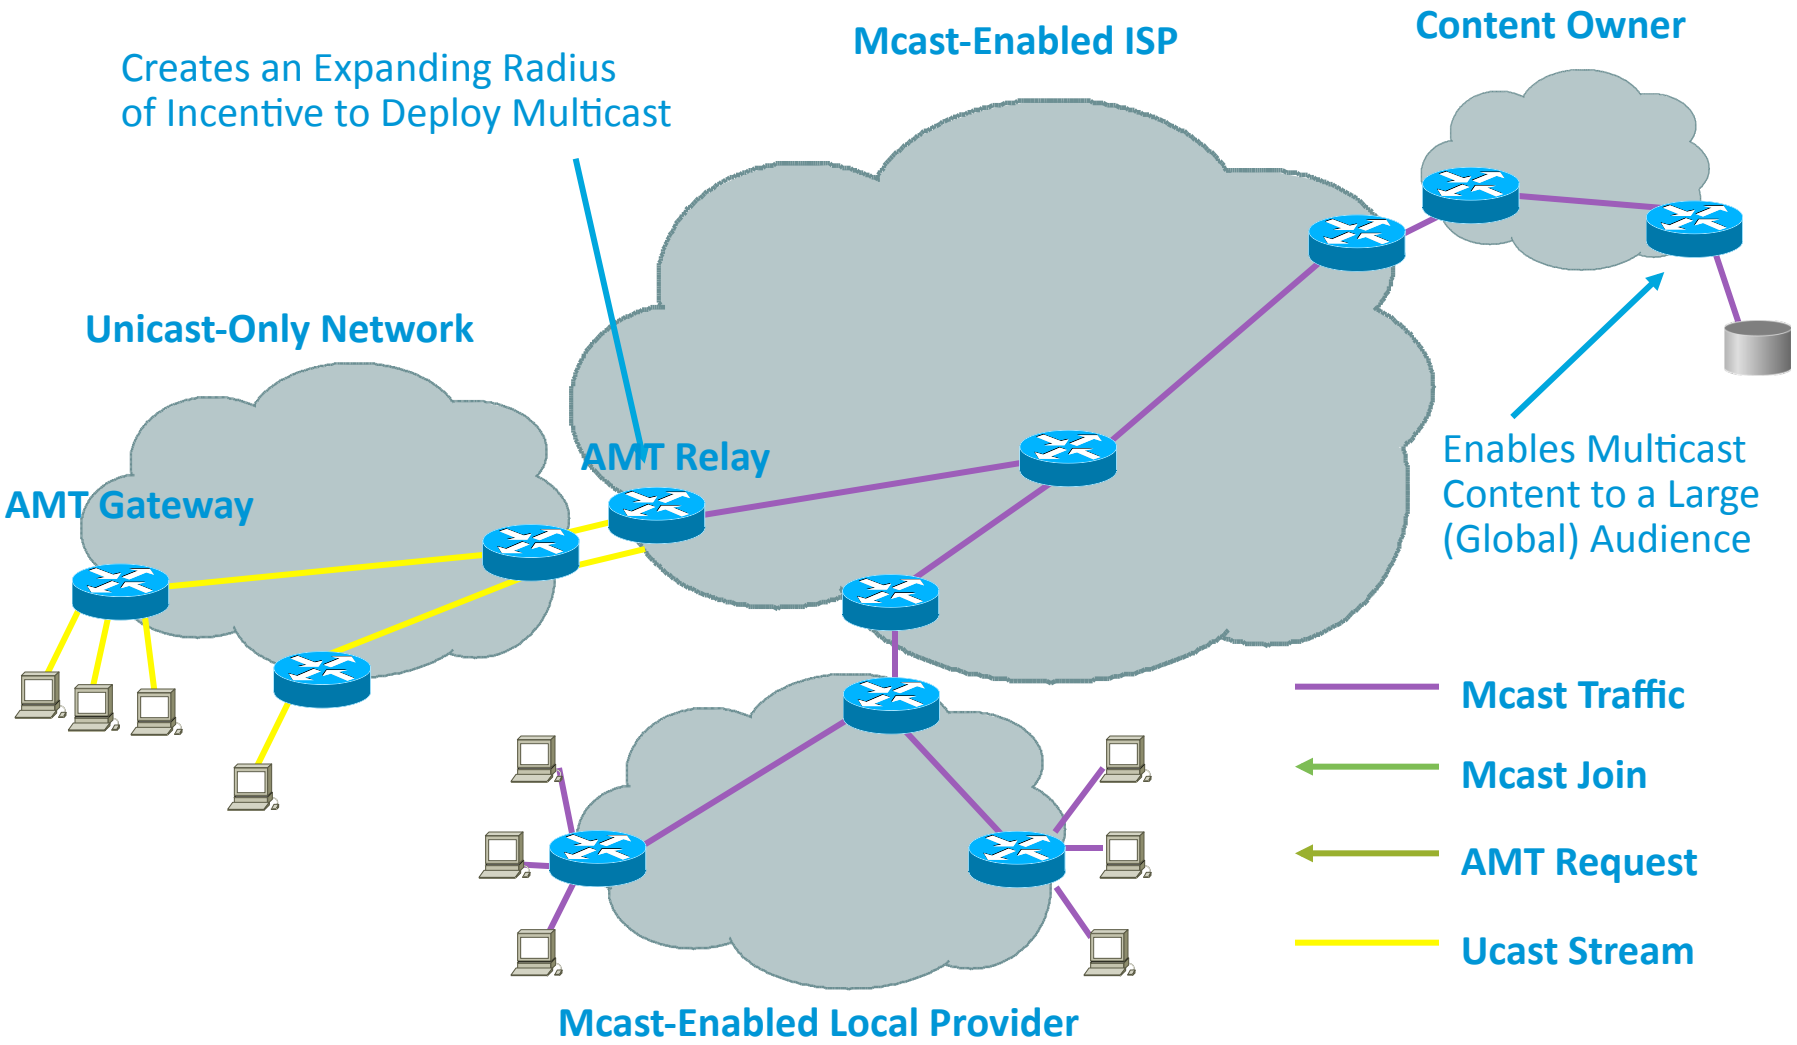

AMT in action – “spreading the word”

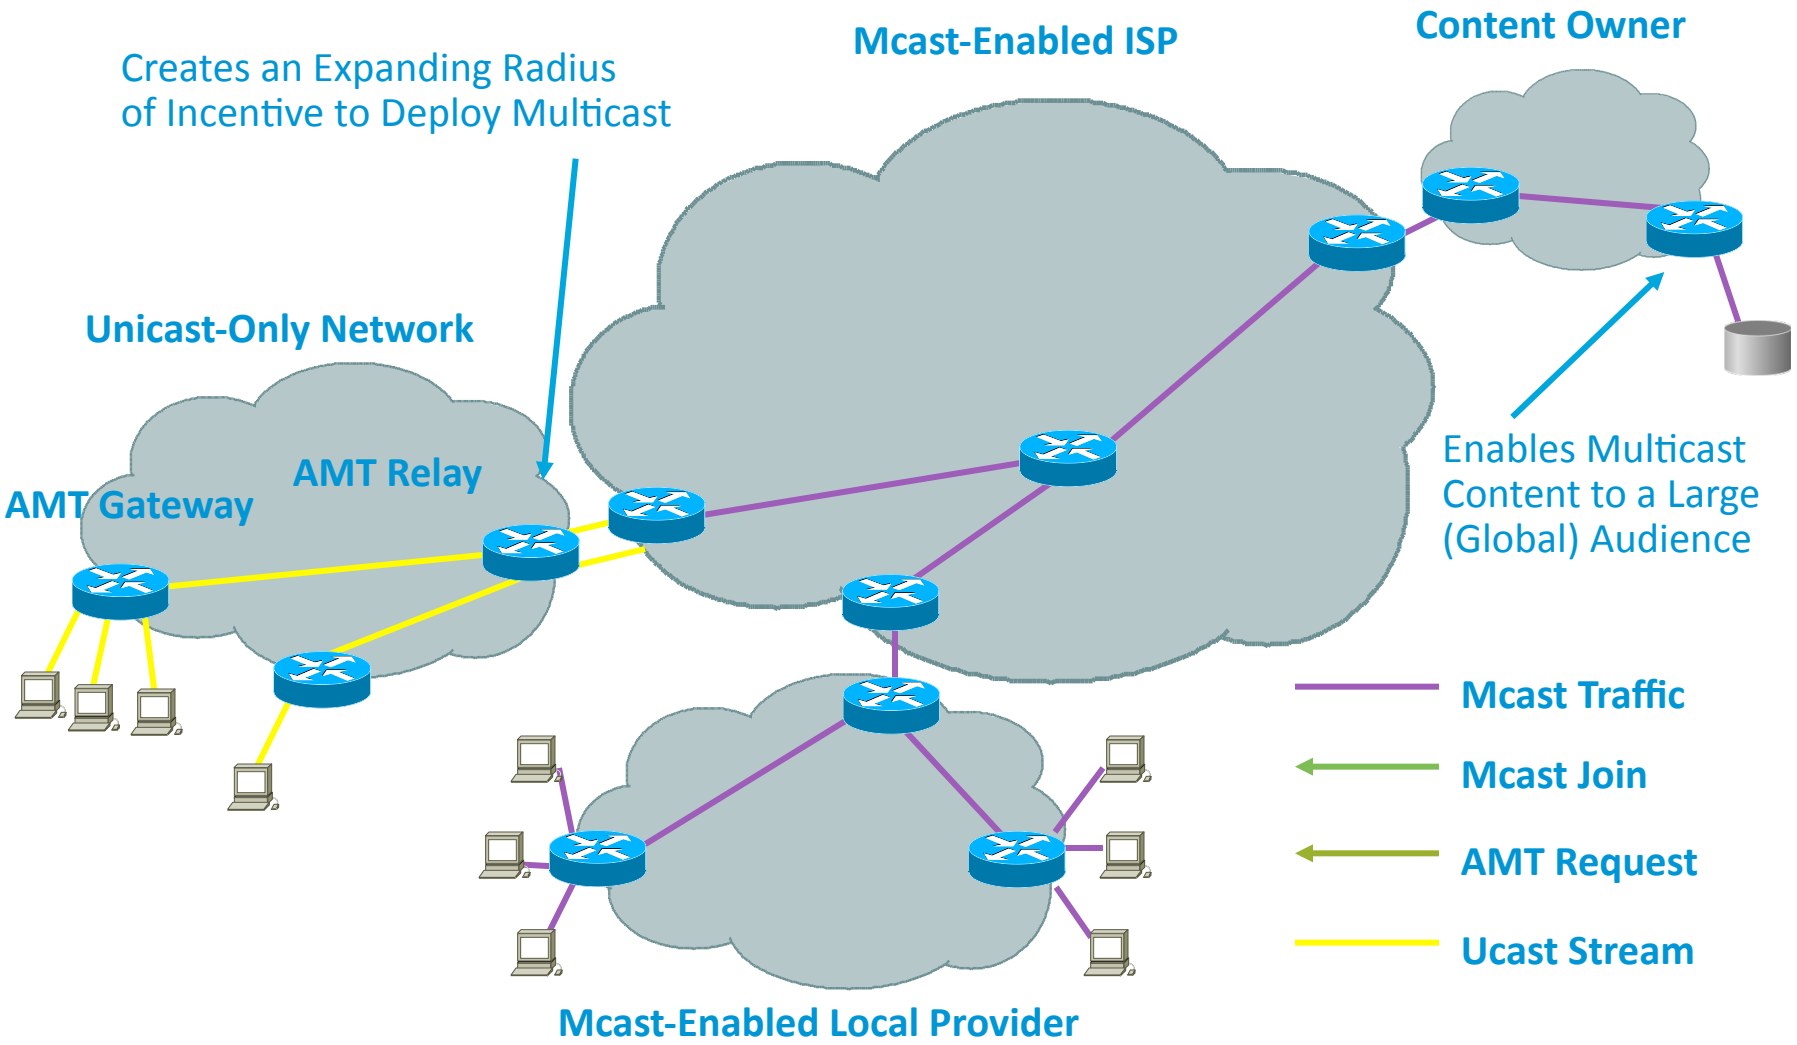

AMT in action – "spreading the word"

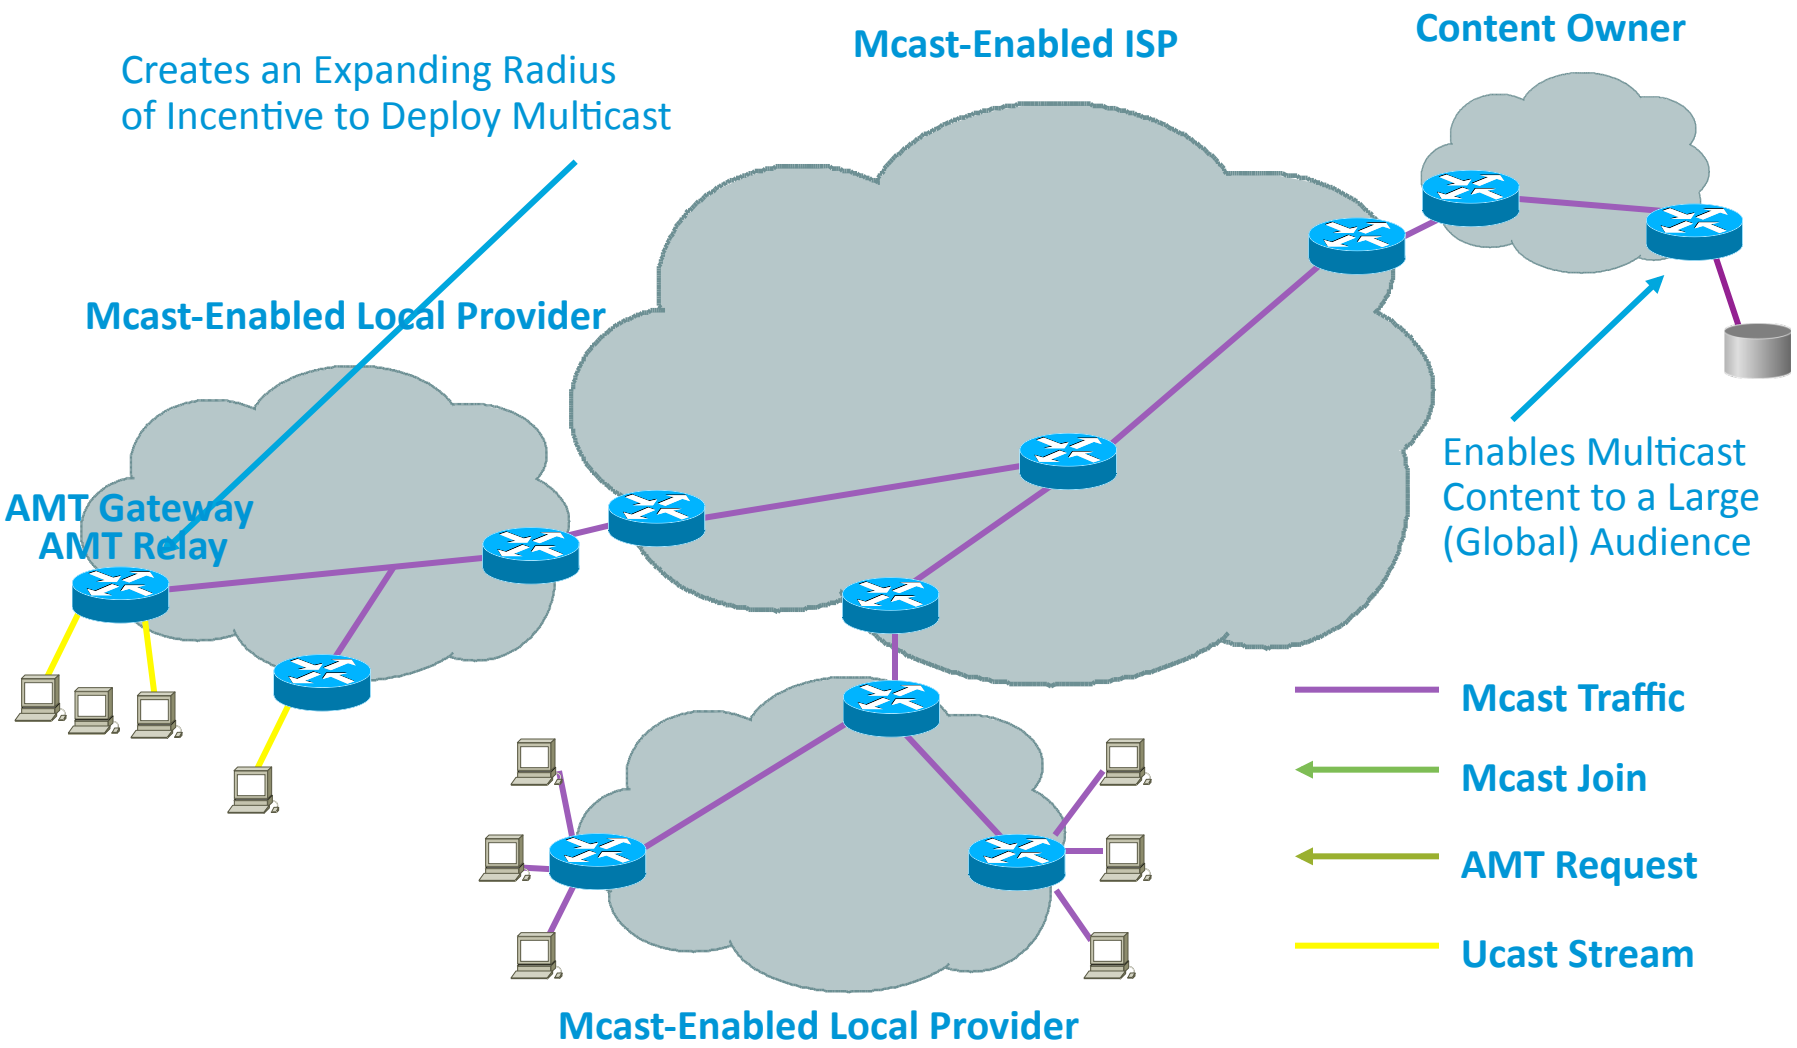

AMT drives global multicast adoption

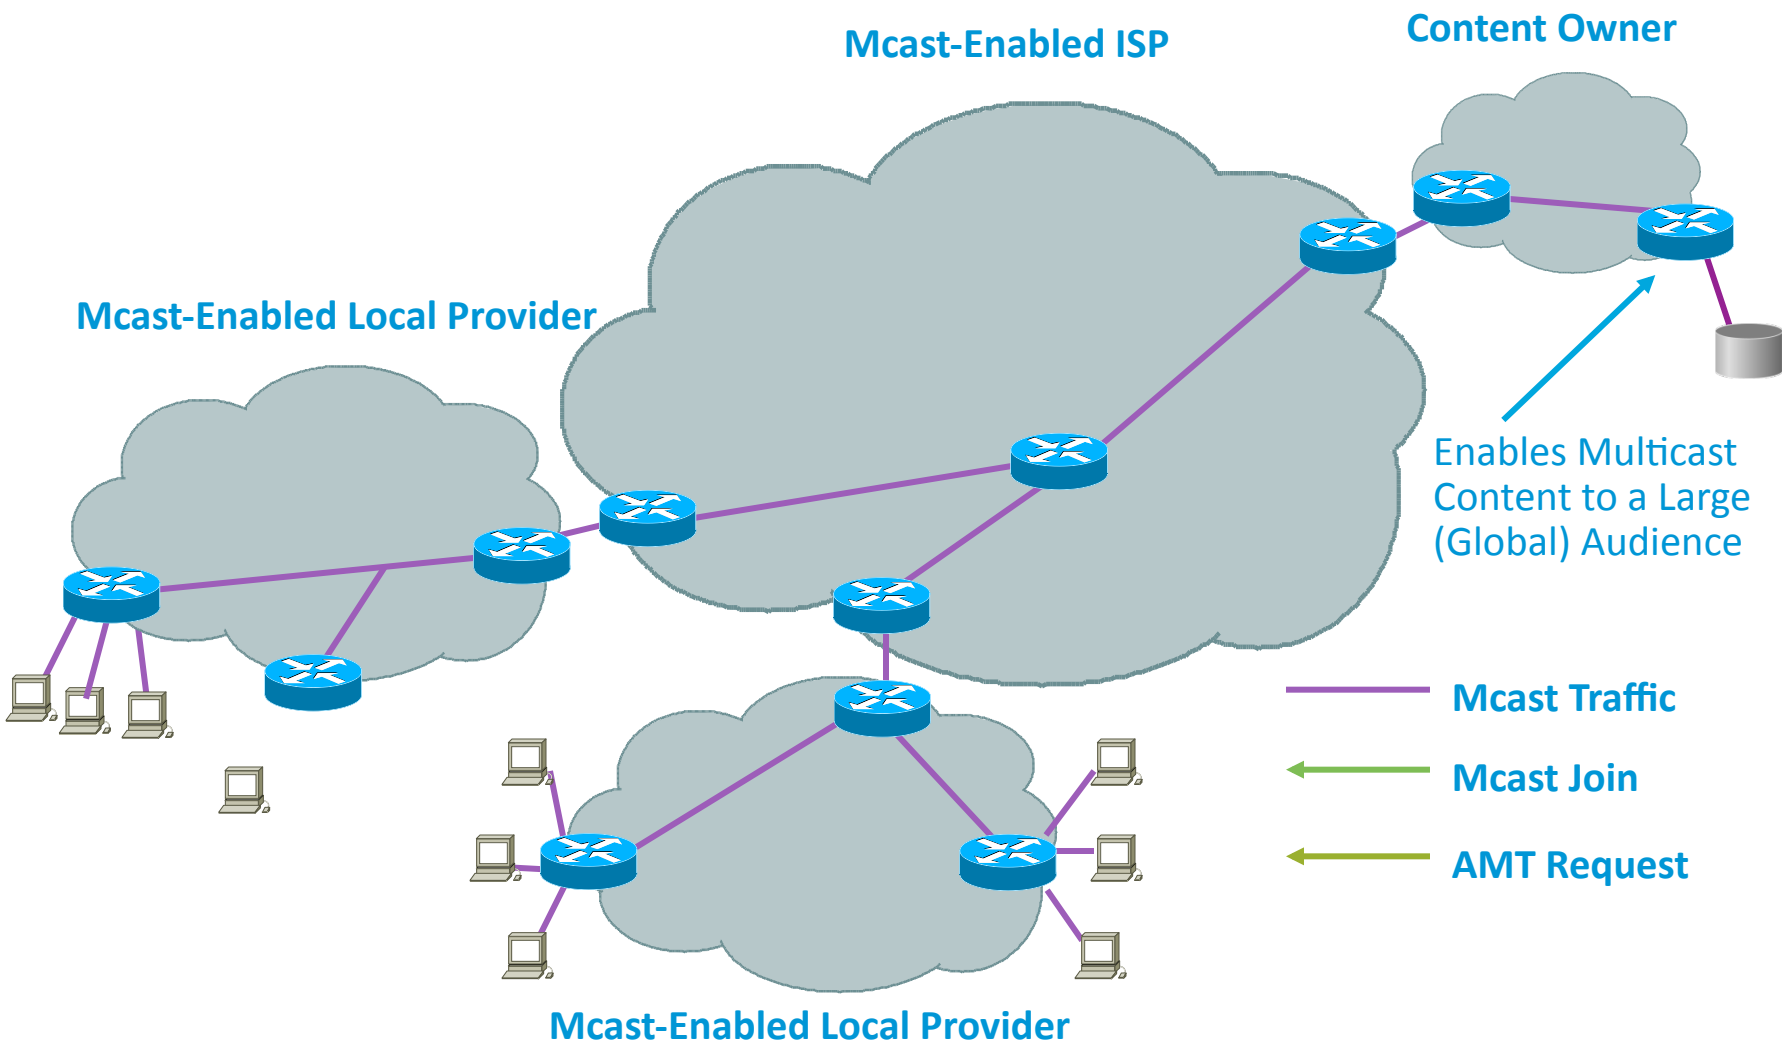

Summary

- How does multicast defer from unicast delivery?
- How are Service Providers deploying multicast in their own access networks for their managed services?
- Can multicast benefit from an IPv6 rollout?
- Why is the Internet not multicast enabled everywhere?
- What solution exists to bridge the multicast islands across the Internet to enable end to end multicast delivery?

AMT References

- Automatic IP Multicast Tunneling:

  http://tools.ietf.org/id/draft-ietf-mboned-auto-multicast

- UTDallas public Relay/Gateway available

  http://cs.utdallas.edu/amt/

- AMT/RTSP Proxy Server:

  http://sites.google.com/site/amtproxy/

- Octoshape AMT enabled client:

  http://www.octoshape.com/?page=showcase/multicast

- Wireshark Dissector for AMT:

  http://www.larkwoodlabs.com/wireshark-dissector-for-amt

- EBU Technical Review (December 2010):

  http://tech.ebu.ch/webdav/site/tech/shared/techreview/trev_2010-Q4_AMT_Kernen_Simlo.pdf

Thank you.