Multicast

Tuomo Karhapää
tuomo.karhapaa@otaverkko.fi
Otaverkko Oy

Agenda

- Why multicast?
- The concept of multicast
- Multicast groups
- Multicast addressing
- Multicast routing protocols
- MBONE
- Multicast applications
- Conclusions

 2

Why multicast?

- Unicast vs. multicast
- typical multicast applications are
  - audio and
  - video transmission
  - collaborating
- multicasting reduces network bandwidth usage

 3

Why multicast?

- example of bandwidth usage:
  - MPEG1 videostream 1.5 Mbps
  - 10 receivers
- bandwidth needed at server site with
  - unicast 15 Mbps
  - multicast 1.5 Mbps

 4

The concept of multicast

- IPv4 supports three types of addresses
  - unicast (point-to-point communication)
  - broadcast (send a packet to entire subnet)
  - multicast (delivery of packets to a set of hosts that have joined to multicast group)

 5

The concept of multicast

- multicast is not connection oriented
- multicast datagram is delivered to destination group members as a unicast IP datagram (UDP)
- unicast IP uses class A, B and C addresses
- multicast IP uses class D address format (224.0.0.0.-239.255.255.255)

 6

The concept of multicast

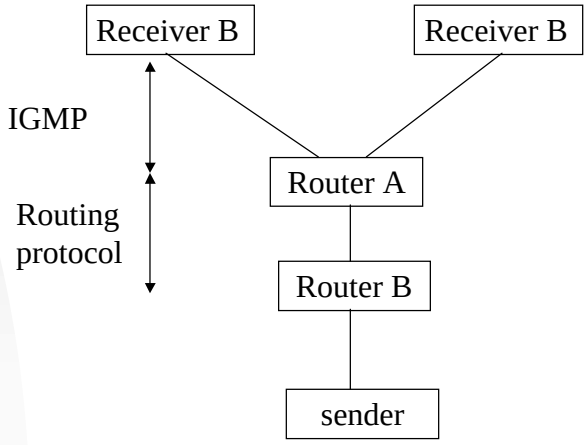

 7

Multicast groups

- hosts are free to join and leave multicast groups
- there are no restrictions on physical locations
- host may be a member of one or more groups
- sender need not to be a member of the group

 8

Multicast addressing

- the format of a 32-bit class D address

1	1	1	0	Multicast Group ID (28 bits)

- IANA maintains a list of registered IP multicast groups
- the base address 224.0.0.0 is reserved

 9

Multicast addressing

- 224.0.0.1-224.0.0.255 is reserved for routing, discovery and maintanance protocols
  - 224.0.0.1 all systems on this subnet
  - 224.0.0.2 all routers on this subnet
  - 224.0.0.4 all DVMRP routers
  - 224.0.0.5 all OSPF routers

 10

Multicast addressing

- routers should not forward a multicast datagram with a destination address in this range
- groups from 224.0.1.0 to 239.255.255.255 are assigned to various multicast applications or remain unassigned
- 239.0.0.0-239.255.255.255 are reserved for site-local applications, not for Internet wide applications

 11

Multicast addressing

- when sender and receivers are in the same subnet, transmission and reception of multicast frames are simple processes
- sender transmits the IP packet to the multicast group, the NIC maps class D address to network level multicast address

 12

Multicast addressing

- when sender and receivers are located to different subnets, things are more complicated
- routers are required to implement a multicast routing protocol for
  - discover delivery tree
  - support multicast data packet forwarding
- each router needs also implement a group membership protocol to learn group members on its directly attached subnetworks

 13

IGMP version 1

- Internet Group Management Protocol (RFC-1112)
- IGMP is used to register a host to multicast group
- routers periodically transmit Host Membership Query messages to determine which groups have members on directly attached networks
- query messages are sent to all-host group (224.0.0.1) with TTL=1

 14

IGMP version 1

- when host receives a Query message it responds with Host Membership report to inform which groups it belongs
- routers need not to maintain detailed list which host belong to each multicast group
- routers only need to know that there is at least one group member

 15

IGMP version 2

- IGMP 2 (RFC-2236) is enhanced version of IGMP 1
- backward compatible with IGMP version 1
- Group-Specific Query: Query is sent to specific multicast group rather than all groups
- Leave Group message

 16

IGMP version 3

- draft specification
- Group Source Message; host can elect to receive traffic from specific sources of a multicast group
- that will help to reduce bandwidth usage because multicast routing protocols use that information when constructing the multicast delivery trees

 17

Multicast forwarding algorithms

- IGMP is used for multicast packet delivery from local router to directly connected subnetworks
- for Internet-wide multicast service there is a need to use multicast routing protocols
- multicast routing protocol is responsible for the constructing delivery trees and multicast packet forwarding

 18

Multicast forwarding algorithms

- there are different algorithms used by multicast protocols
  - Flooding
  - Spanning trees
  - Reverse Path Broadcasting (RPB)
  - Truncated Reverse Path Broadcasting (TRPB)
  - Reverse Path Multicasting (RPM)
  - Core-Based Trees

 19

Flooding

- simplest technique for delivering multicast datagrams
- if the packet is seen first time, router forwards packet to all interfaces
- if router has seen the packet before, it discards packet
- flooding is very simple to implement
- it generates duplicate packets and uses resources inefficiently

 20

Spanning trees

- more efficient method than flooding
- spanning tree is a structure of active paths
- packets are forwarded to active paths except to the path which originated the packet
- powerful method and quite easy to implement

 21

Reverse Path Broadcasting (RPB)

- source-rooted spanning tree
- router forwards packet to all interfaces (child) except incoming interface (parent)
- distance-vector routing protocol's routing table may be used to determine shortest path back to the source
- need not to know hole tree structure

 22

Truncated Reverse Path Broadcasting (TRPB)

- RPB forwards packets to all leafs
- TRPB eliminates unnecessary traffic on leaf subnetworks by determining group membership via IGMP
- but it does not consider group membership when building the distribution tree

 23

Reverse Path Multicasting (RPM)

- enhancement of RPB and TRPB
- distribution tree consist of
  - subnetworks with group members
  - shortest path to subnetworks with group members
- prune mechanism; packets are forwarded only to branches that lead to members of group

 24

Reverse Path Multicasting (RPM)

- not scalable; all routers need to maintain state information of all groups and sources
- multicast packets must be periodically forwarded to every router in the internetwork to maintain prune information

 25

Core-Based Trees (CBT)

- RFC-2189
- source-rooted, shortest-path delivery tree for each source, group pair
- similar to the spanning tree algorithm except it has different core-based tree for each group
- it only requires a router to maintain state information for each group, not for group, source pair

 26

Core-Based Trees (CBT)

- conserve bandwidth since it does not to forward multicast packets peridically to every router
- the bottleneck is near a core router because traffic from all sources traverses via same links

 27

MOSPF

- RFC-1584
- multicast extensions to OSPF version 2
- MOSPF is defined in RFC-1584
- uses network topology constructed by unicst OSPF
- do not support tunnels

 28

MBGP

- Multiprotocol Extensions for BGP-4 (RFC-2283)
- supports different unicast and multicast routing policies concurrently
- scalable

 29

DVMRP

- Distance Vector Multicast Routing Protocol (RFC-1075)
- uses Reverse Path Multicasting protocol
- first implementation was mrouted
- two types of interfaces
  - physical interface to a directly connected subnetworks
  - tunnel interface to another multicast island

 30

DVMRP

- all interfaces are configured with metric and TTL threshold

Initial TTL	Scope
0	Restricted to the same host
1	Restricted to the same subnet
32	Restricted to same site
64	Restricted to same region
128	Restricted to same continent
255	Unrestricted in scope

 31

DVMRP

- multicast packet is forwarded if the TTL value in the IP header is greater than the TTL threshold assigned to the interface
- the initial packet is forwarded to all interfaces to determine prune information
- if there are no members on certain leaf, that branch is removed from delivery tree

 32

DVMRP

- if pruned branch discovers new group member it sends graft message to upstream to cancel prune state
- upstream router restored that branch to delivery tree
- DVMRP maintains routing and forwarding tables with integrated routing protocol

 33

PIM

- Protocol Independent Multicasting
- PIM is independent on the mechanism provided by any particular unicast routing protocol
- it requires some unicast routing protocol to determine network topology and topology changes
- two modes: dense and sparse

 34

PIM-DM

- PIM-DM uses Reverse Path Multicasting
- quite similar to DVMRP
- uses Reverse Path Multicasting (RPM)
- PIM-DM is developed for campus LANs where group membership is usually dense and there is also more bandwidth available than WANs

 35

PIM-DM

- differences between PIM and DVMRP
  - uses existing unicast routing protocol
  - simply forwards all multicast packets to downstream until prune message is arrived

 36

PIM-SM

- RFC-2362
- PIM-SM provides a efficient multicast routing mechanism for communication between sparsely distributed groups
- solves dense mode protocols' scaling problems by limiting multicast packets to routers which are interested in receiving traffic

 37

PIM-SM

- PIM-SM differs from existing dense-mode algorithms by following ways
  - routers are required to join to sparse-mode distribution tree by transmitting join messages
  - if router is not part of that distribution tree, it will not receive multicast traffic

 38

PIM-SM

- PIM-SM uses concept of Rendezvous Points (RP)
- at the RP receivers meet sources
- the initiator of group selects the primary RP and alternative RPs
- only one RP is active
- the host which wants to join to the group sends join message to local router which sends explicit join message to group's primary RP

 39

MBONE

- Multicast backbone on the Internet
- first transmission was IETF meeting "audiocast" 1992
- MBONE is a virtual network on the top of Internet
- used for audio and video transmissions and for collaborating
- small part of Internet's backbone routers supports multicasting

 40

MBONE

- typical sessions are
  - IETF meetings
  - radio stations
  - NASA shuttle missions
  - events like rock concerts
- typical bandwidth usage
  - for audio 13 kbps - 64 kbps
  - for video up to 128 kbps

 41

MBONE

- mostly used by universities and research institutes
- most of the commercial operators do not support multicasting
- problems with firewalls, one solution is a tunnel interface
- usually low bandwidth transmission

 42

Multicast applications

- at the moment different UNIX systems and also Windows OS support multicasting
- main router vendors support multicasting
- there is free client software available
- the “killer” application is still missing

 43

Multicast applications

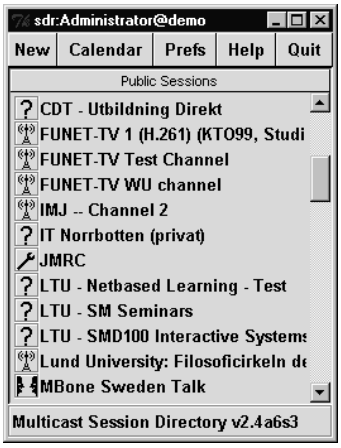

 44

Multicast applications

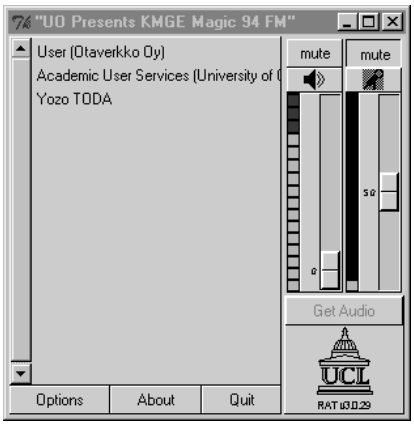

 45

Multicast applications

- MP3 multicasting
  - offers reasonable quality for audio streaming
  - RTP used for payload format
  - might be “killer application”

 46

Conclusions

- multicasting reduces bandwidth usage
- not easy to implement
- needs protocol development work for Internet-wide usage

 47

More information

- www.ietf.org
- www.mbone.com

 48