Multicast as an ISP service

Lecture slides for S-38.3192
16.2.2006
Mika Ilvesmäki

Goals of this lecture

- After this lecture you will be able to
  - Give an overall technical view of multicast
    - Addressing, protocols, routing
    - Please note, that Multicast (as a technology) is also lectured at least in 38.2122 and 38.3188
  - List different application and service types that could be deployed with multicast
  - Discuss the problems in bringing Multicast to production use in ISP network
  - Discuss the different aspects of integrating QoS solutions to Multicast

Packet distribution I

- TCP/IP stack implements four types of packet distribution methods:
  - **Unicast**: From one to another host
    - Normal point-to-point IP traffic
    - The packet is sent only to the receiver
  - **Broadcast**: From one host to all hosts on a network
    - "Is there anybody out there?"
    - The packet is seen by all hosts on the (local) network

Packet distribution II

  - **Anycast**: From one host to one host in a group of hosts
    - To access a service served by multiple hosts (in IPv6)
    - The packet is received by one of the receivers
  - **Multicast**: From one to a group of hosts
    - All group members receive the packet
    - Packet duplicated only when necessary

Many-to-many traffic

- How to send "many-to-many" traffic?
  1. Central servers: Servers duplicate packets to all members of the group (Unicast).
  2. Connected mesh: Every group member has a connection to all or some others (Unicast/Broadcast).
  3. Multicast: Network duplicates the packets as needed.
- First two choices do not scale and waste resources (bandwidth, especially).

Multicast addressing in IPv4

- Multicast traffic uses class D IP-addresses:
  - Flat hierarchy
  - 224.0.0.0 - 239.255.255.255
    - 224.0.0.0 – 224.0.0.255 reserved for multicast protocols
    - 224.0.1.0 – 224.255.255.255 reserved for various applications
    - 225.0.0.0-238.255.255.255 for free use
    - 239.255.0.0 – 239.255.255.255 for private use (administratively scoped)
      - Not just TTL-bounded, limited vendor support
- See http://www.iana.org/assignments/multicast-addresses for complete list.

Multicast addressing in IPv6

- Address structure

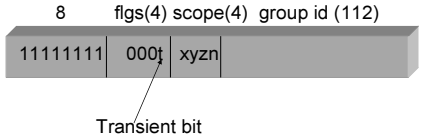

- The scope field determines the maximum "depth" of the multicast group
  - 1=Node local, 2=Link local , 5=Site local, 8=Organization local, E_{16}=Global scope

Multicast routing

- Routers form multicast trees
  - The tree leaves are the receivers
  - The trees are maintained with multicast routing protocols and often referred to as spanning trees
  - Trees come in several formats
    - Shortest Path from each source
    - Shortest path from a RP (core of the tree)
    - Unidirectional shared tree
    - Bi-directional shared trees
- Multicast routing protocols include
  - DVMRP, MOSPF, PIM-SM, PIM-DM, CBT, OCBT, HIP and BGMP, to name a few.

Multicast packet forwarding

- RPF-check in the router:
  - All multicast packets are checked for the source address (not for the destination as in unicast forwarding)
    - The incoming interface must be an interface pointing towards the source address (sender)
      - If this checks, packet is forwarded to appropriate interfaces, otherwise the packet is dropped.

Multicast group management

- A host must *join* a specific group to receive the traffic in that group but can send to a group without joining.
  - Non-member senders!
- Membership is controlled by the *IGMP* protocol.
  - IGMPv2 is the most commonly used (and an IETF standard)

IGMPv2

- Backwards compatible with IGMPv1
- 3 message types
  - Membership Query (general or group specific)
  - Membership report (general (v1) or addressed to the group (v2))
  - Leave Group
- IGMPv3 adds support (to Membership Query msgs) for a system to report interest in receiving packets *only* from specific source addresses, or from *all but* specific source addresses, sent to a particular multicast address.

Hardware for ISP

- De facto architecture
  - IGMPv2, DVMRP, MOSPF, PIM-SM (within autonomous systems) and MSDP or MBGP (for interdomain multicast routing)
  - These protocols are readily usable with most commercial routers
    - However, most of the time these are turned off!
    - Why?

Why Multicast?

- Multicast saves bandwidth
  - Multicast is desirable for sources(servers) and administrators
    - Sources may scale service to large number of users
    - Redundant connections could be joined
  - Receivers do not care whether content is received via unicast or multicast
    - BW/QoS/realtime requirement is the same anyway

Application drivers

- Unicast killer apps: Email+WWW
- Multicast apps (one-to-many or few-to-few)
  - Audio & Video distribution
  - Push applications (PointCast)
  - Audio & Video conferencing
  - File transfer (updates, file sharing, web caching, distributed databases)

MBone multicast service

- Virtual network (using the Internet unicast resources) transmitting multicast-packets
  - Tunneling
- Multicast is a method of routing and delivering packets in a many-to-many situation
- MBone is one implementation of multicast method.

Multicast in MPLS

- Manually with explicit trees
  - Requires extra work, but is an option in small networks
- MPLS networks utilize Multicast Distribution Trees (MDT)
  - Point-to-multipoint and multipoint-to-multipoint LSPs
- MPLS/BGP/VPNs
  - In unicast VPNs the P routers do not maintain per-VPN state (only per-PE)
  - In multicast VPNs the P routers have to maintain some information on active multicast trees (to avoid forwarding unneeded traffic).
- MPLS+Multicast is still very much work in progress

IP multicast as a service

- Open service –model
  - No restrictions on creating multicast groups, receiving from a group or sending to a group
  - Group membership means that one is reachable to other members of the group
    - No access control
  - Multiple senders may share the same multicast address
    - No pre-emption or address reservation
  - Group size is dynamic and unknown, member status is unknown
- Summary: IP multicast group is not managed (or the management is distributed at best)

Customer requirements

- ISP is a customer-driven business
- Any service offered must add value to the customer
- For multicast to add value, customers must have
  - Global access to the service
  - Ease and transparency of install and use
  - Controlled group membership
  - Unique addressing (for the session)
  - Reliable transmission

Management

- Lack of NAT support (work in progress)
- Firewalls do not recognize multicast addresses
  - Solution: tunnels. Consequence: A hole in the wall. Another solution: static routes (and more human management and therefore more expensive service)

Domain independence

- Traffic sources in other domains are out of the ISP control
  - Traffic control
  - Service level
  - Transit traffic is not preferred
  - Control of address advertisements

Multicast benefits for the ISP

- Multicast reduces BW required to transport data to multiple recipients
  - Multicast is expensive (aka H-A-R-D) to deployment, installation and management
  - Are BW savings higher than cost of deployment and management?

Requirements for IP multicast service

- All routers must support multicast (Customers want global access)
- Domain independence
- Additional requirements
  - Group management
    - Authorization for group creation, receiver and sender authorization
  - Distributed multicast address allocation
  - Security
    - Protection against attacks on multicast routes and sessions
    - Support for data integrity mechanisms.
  - Support for network management
  - Billing Multicast

Multicast and QoS I

- RSVP + IntServ
  - IntServ built with multicast in mind
    - Receiver requests QoS based on Sender description of the service
    - Unidirectional reservations
    - Filters (Shared Explicit, Wildcard, Fixed Filter) illustrate different types of multicast sessions
    - Reservation merging
  - Usual IntServ problems persist
    - Scalability issues
    - Local management (CAC, policy control)

Multicast and QoS II

- DiffServ
  - There should be no per-group state in the core -> reservation uncertainty
  - Multicast group size not defined in advance and DiffServ not having per-user access control -> unknown resource usage
  - Replication of packets in the network may cause problems in capacity handling
  - Problem of relative QoS
    - PDB in one domain is not necessary the same PDB in other domains -> need for dynamic DSCP allocation -> not a standardized approach

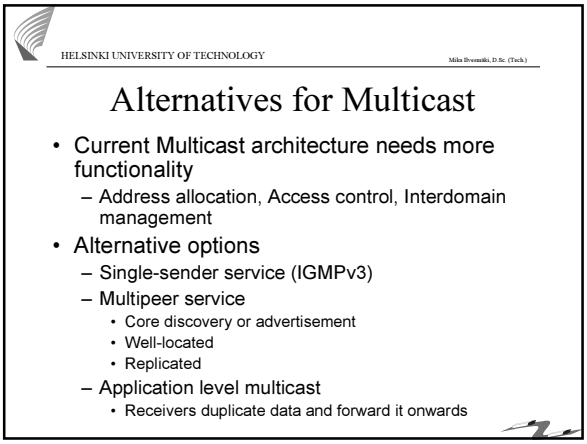
HELSINKI UNIVERSITY OF TECHNOLOGY
Miska Ilvesmäki, D.Sc. (Tech.)
Alternatives for Multicast
• Current Multicast architecture needs more functionality
– Address allocation, Access control, Interdomain management
• Alternative options
– Single-sender service (IGMPv3)
– Multipeer service
• Core discovery or advertisement
• Well-located
• Replicated
– Application level multicast
• Receivers duplicate data and forward it onwards