RTP Topologies

Clue Interim June 2012

Magnus Westerlund
Bo Burman
Ericsson

Outline

- The Big Issue
- Goal
- Evaluation Criteria
- Topologies
- Conclusions

The Big Issue

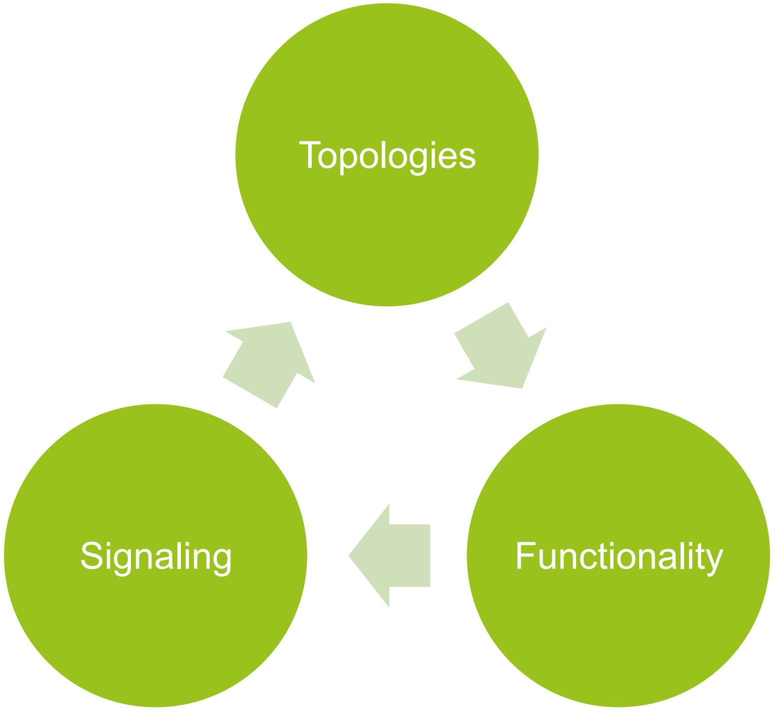

Signaling can restrict what topologies that are supported, and thus the functionality a CLUE system may have!

Goal

- Present what **RTP functionality** a given topology enables
- Start a discussion on what topologies and media plane functions CLUE wants to support
  - To ensure correct requirements on the signaling

Evaluation Criteria

- Security
  - Key-management
  - Nodes that are in the security context (who have the keys)?
    - Trust Structures
  - Source Authentication
    - End-to-End verifiable
    - Trust in central node

Evaluation Criteria

- Congestion Control
  - Multi-hop
    - Need for information and requests to bridge across hops
  - One or Multiple Receivers of the same RTP stream
    - Meet requirements from multiple end-points
  - Transcoding
    - Enabling bit-rate adjustments
    - Breaking multi-hop control loops
  - Media Aggregate adjustments
    - Prioritization between streams

Evaluation Criteria

- Source Identification
  - Receiver must be able to determine source of media
    - Reference in Meta-Information
    - Identity for Control Requests
  - Media mixing or compositions
    - Multiple contributing sources
  - Translation of source identification information
    - Require additional layer of identification labels
    - OR
    - Force all end-point communication through node that translates
  - Conference Wide common identity space required?

Evaluation Criteria

- Bandwidth Consumption
  - Deliver most appropriate media properties
    - Transcode
    - Choice from Simulcast alternatives
    - Source Codec Control
  - Select the N out of M streams most needed by the application
    - Possibility to Prune unneeded streams
  - Mix or composite N streams into one
  - Translation transcode / re-encode
    - Increases bandwidth usage to maintain quality

Evaluation Criteria

- Media Quality
  - Transcoding / re-encoding
    - Quality reduced per spent bit
  - Delay
    - Need to be kept low in to maintain interactivity
    - Inter continental communication
  - The rate vs distortion relation is approximately logarithmic
    - Quality gain per bit will affect prioritization between:
      - Increasing one stream's quality
      - Allowing additional streams
  - Difficult trade-off between quality, delay, bit-rate consumption and functionality

Evaluation Criteria

- Distribution of Complexity
  - Various factors, e.g.
    - Processing
    - Memory
    - Implementation cost
  - Depending on Topology
  - Some complexity can be moved between central nodes and end-points
    - Impact on a central node can be different from an end-point for a given functionality
  - Node Limitations must be taken into account
    - Forces location of functionality
    - Can cost quality

Topologies Outline

- Point to Point
- Distributed End-point
- Multi-Unicast (MESH)
- Mixers
  - Media Mixer
  - Media Switching
  - Source Projection
- Relay (Transport Translator)
- Selective Forwarding
- End-point Forwarding
- Any Source Multicast
- Sender Source Multicast

Point to Point

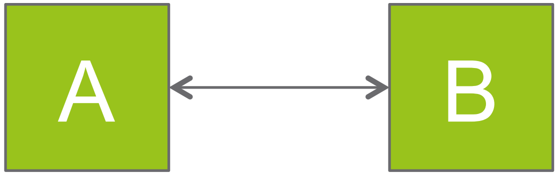

- Description:
  - One peer communicate directly with a single peer over unicast
- Security:
  - Authentication of Peer / User
  - Confidentiality and Integrity between peers
- Congestion:
  - Receiver can report statistics or request direct adaptation from sender
- Identification:
  - Senders sources map one to one with RTP media streams
- Bandwidth:
  - Receiver can request media to be tailored to its needs
  - Action to increase or decrease bandwidth can depend on the current path capacity
    - At high capacity add additional streams to provide additional functions
    - At low reduce to single stream and focus at maximize quality for the most important content
- Quality:
  - Optimal in relation to Path Capacity and Properties
- Complexity:
  - All in the end-points
  - Limitations in end-point directly affect what sender can produce and receiver can accept

Distributed End-point

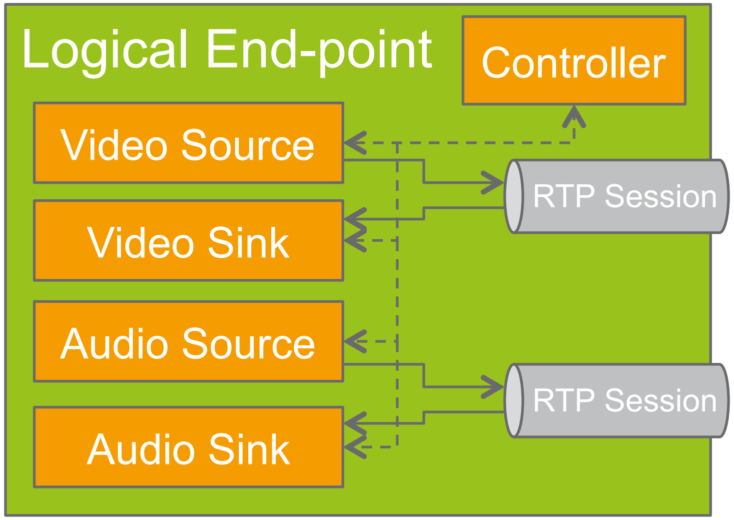

- Description:
  - Distributed realization of a logical end-point
  - Different IP addresses for various components
    - Camera (Video source)
    - Display (Video sink)
    - Microphone / Audio mixer (Audio source)
    - Loudspeaker (Audio sink)
    - Controller (Signaling end-point)
  - There can be multiple instances of one component type
- Security:
  - **Each source or Sink must be keyed with the other end-points key(s)**
  - Controller responsible to provide logical end-point identity
- Congestion:
  - Receiver component can report statistics or request direct adaptation from media sending component
  - **Prioritization between media streams in the aggregate are complicated by distribution**
  - Due to different source / destination addresses network load balancer may give different routes to different flows
- Identification:
  - Senders sources map one to one with RTP media streams
  - A logical end-point may have multiple presences in an RTP session due separation of sources and sinks
  - **Multiple different IP addresses or hidden behind aggregation point**
- Bandwidth:
  - **Trade-off between centralized control and distributed handling of adaptation and prioritization**
- Quality:
  - May become sub-optimal in relation to Path Capacity and Properties due to control latencies
- Complexity:
  - Additional complexities for control within the end-point

Multi-Unicast (Mesh)

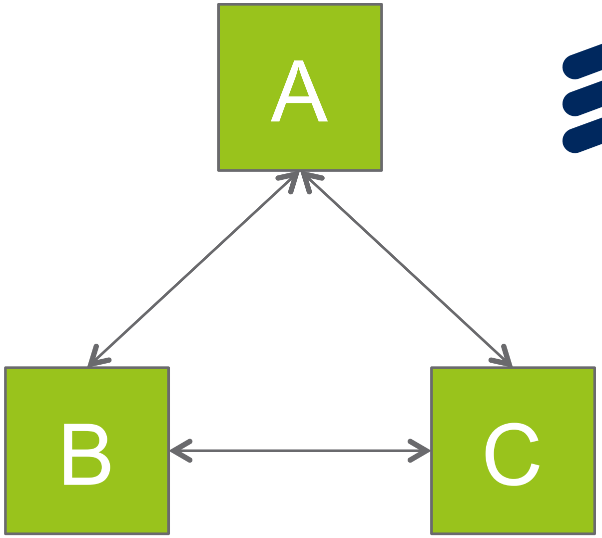

- Description:
  - One peer communicate directly with multiple peers
  - Each peer to peer communication is independent unicast
  - **Each peer pair can have its own RTP session**
- Security:
  - Individual Authentication of each Peer / User
  - Confidentiality and Integrity between pair of peers
- Congestion:
  - Receiver can report statistics or request direct adaptation from sender
  - **All Peers will commonly share first hop/hops and the available capacity / bottleneck**
  - Sender can produce independently encoded media or produce one encoding sent to multiple peers.
- Identification:
  - Sender's sources map one to one with RTP media streams within one RTP session
  - **Using multiple RTP sessions results in independent SSRC/CSRC spaces between the sessions**
    - Could select to use unique values over multiple RTP sessions or use different layer
- Bandwidth:
  - **Receiver can request media to be tailored to it's needs**
    - **May be forced to accept a compromise based on other paths in case sender share media encoder**
  - Desirable to enable different trade-offs based on path capacity
- Quality:
  - Can be optimal in relation to Path Capacity and Properties
  - **To reduce sender complexity in encoding less than optimal quality may be received**
- Complexity:
  - All in the end-points
  - Limitations in end-point directly affect what sender can produce and receiver can accept
  - Trade-off in amount of complexity each pair of peers create can affect conference properties

Media Mixer

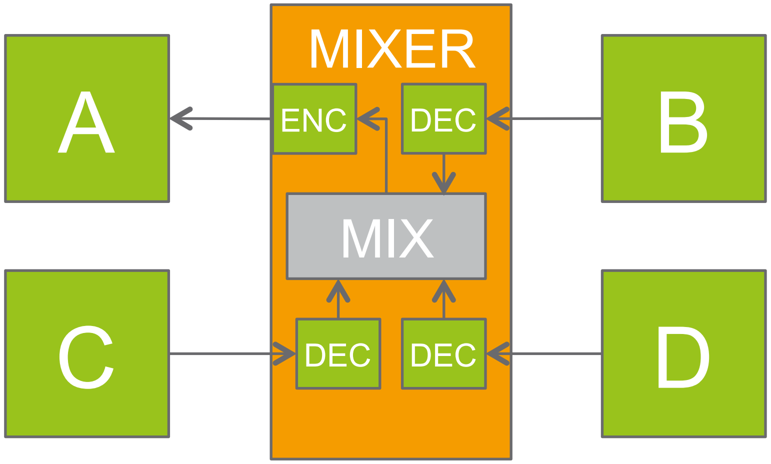

- Description:
  - One peer communicates only with the Mixer
  - Each peer to mixer communication is independent unicast
  - **Mixer provides a mixed or composited media source based on the media streams from the other participants**
  - Each communication can have its own RTP session, or Mixer can create a conference-wide RTP session by sharing SSRC / CSRC
- Security:
  - Mixer handles Authentication of each Peer / User
  - **Mixer is trusted entity and enforcer of some security functions**
  - Confidentiality and Integrity between peer and Mixer
- Congestion:
  - Receiver (Mixer or peer) can report statistics or request adaptation from sender (Mixer or peer) on their path
  - Mixer can choose to forward report / request information (unaltered or aggregated) between paths
  - **Mixer typically produce independently encoded media to each peer, but may re-use some media between receiving peers**
- Identification:
  - **Sender's sources are only visible as contributing sources in Mixer's RTP media streams**
  - Using multiple RTP sessions results in independent SSRC/CSRC spaces between the sessions
    - Could select to use unique values over multiple RTP sessions or use different layer
- Bandwidth:
  - **Mixer can reduce the number of concurrent media streams to a single per media type**
  - Receiver (also Mixer) can request media to be tailored to its needs
- Quality:
  - Maximum Quality limited by participant to mixer path capacity
  - **Quality loss and delay increase in decoding encoding cycle**
- Complexity:
  - **Mixer has one end-point complexity per end-point in the conference, plus media composition and some Mixer-specific logic**
  - Mixer proxies limitations in end-point affecting what sender can produce and receiver can accept, but may add further limits

Media Switching Mixer

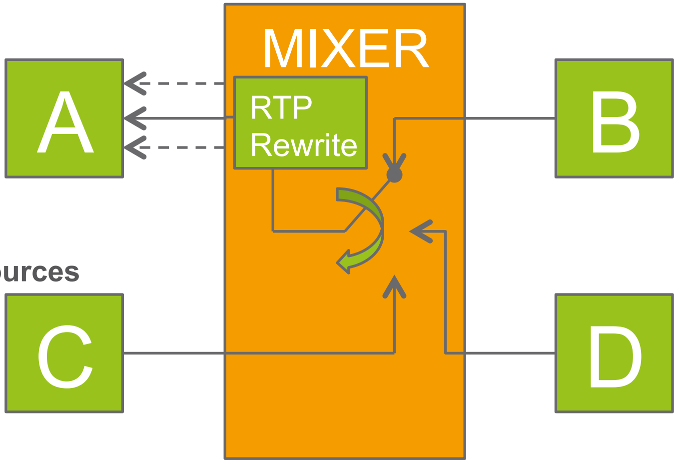

- Description:
  - One peer communicate only with Mixer
  - Each peer to mixer communication is unicast with mixer feedback
  - **Mixer provides one or more conceptual sources selecting original sources**
  - Mixer creates a conference-wide RTP session by sharing SSRC / CSRC
- Security:
  - Mixer handles Authentication of each Peer / User
  - Mixer is trusted entity and enforcer of some security functions
  - Confidentiality and Integrity between peer and Mixer
- Congestion:
  - Receiver (Mixer or peer) can report statistics or request adaptation from sender (Mixer or peer) on their path
  - **Mixer needs to aggregate and forward report / request information between paths, based on some policy**
  - **Mixer distributes encoded media to multiple peers, making single receiver limitation affect more receivers**
  - **Mixer can make use of simulcast or scalable media encoders from senders to adapt to a peer**
- Identification:
  - Sender's sources are only visible as contributing sources in Mixer's RTP media streams
  - Mixer can have multiple SSRCs representing different conceptual media sources
- Bandwidth:
  - Receiver (also Mixer) can request media to be tailored to its needs, but will typically also affect other receivers
  - **Desirable to limit the amount of trade-off based on path capacity**
  - Simulcast and scalability can be used to meet different bandwidth needs or requirements
- Quality:
  - Trade-off of end-to-end Path Capacity and Properties between receivers sharing media from the same sender
  - **Avoids transcoding and its quality reduction and delay penalty**
- Complexity:
  - **Mixer has *no* end-point complexity per end-point in the conference, only switching and some Mixer-specific logic**
  - Mixer proxies limitations in end-point affecting what sender can produce and receiver can accept

Source Projection Mixer

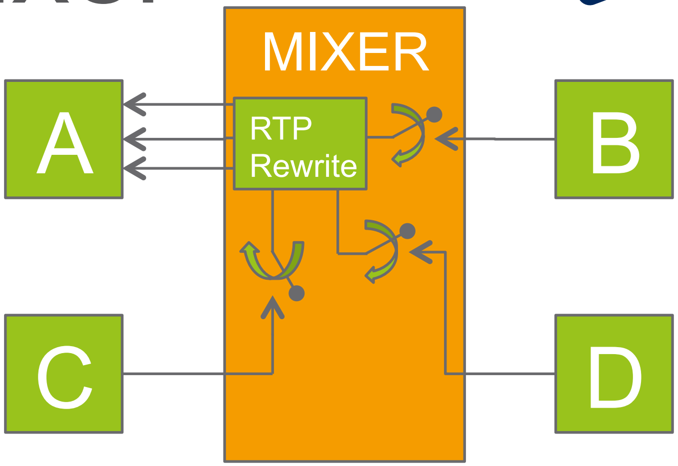

- Description:
  - One peer communicate only with Mixer
  - Each peer to mixer communication is unicast with mixer feedback
  - **Each participant have its own RTP session with Mixer**
  - **Each conference media source is projected into each RTP session**
- Security:
  - Mixer handles Authentication of each Peer / User
  - Mixer is trusted entity and enforcer of some security functions
  - Confidentiality and Integrity between peer and Mixer
- Congestion:
  - Receiver (Mixer or peer) can report statistics or request adaptation from sender (Mixer or peer) on the closest link
  - **Mixer needs to aggregate and forward report / request information between links, based on some policy**
  - **Mixer distributes encoded media to multiple peers, making single receiver limitation affect more receivers**
  - Mixer can make use of simulcast or scalable media encoders from senders to adapt to a peer
- Identification:
  - **Each media source is one to one mapped to a SSRC in Participants RTP session**
  - Sender's SSRC may be renumbered by Mixer, thus requiring RTP-external identification for E2E identity
- Bandwidth:
  - Receiver (also Mixer) can request media to be tailored to its needs, but will typically also affect other receivers
  - Desirable to **limit** the amount of trade-off based on path capacity
  - Simulcast and scalable encoding can be used to meet different bandwidth needs or requirements
- Quality:
  - Trade-off of end-to-end Path Capacity and Properties between receivers sharing media from the same sender
  - **Avoids transcoding and its quality reduction and delay penalty**
- Complexity:
  - **Mixer has *no* end-point complexity per end-point in the conference, only switching and some Mixer-specific logic**
  - Mixer proxies limitations in end-point affecting what sender can produce and receiver can accept

Relay (Transport Translator)

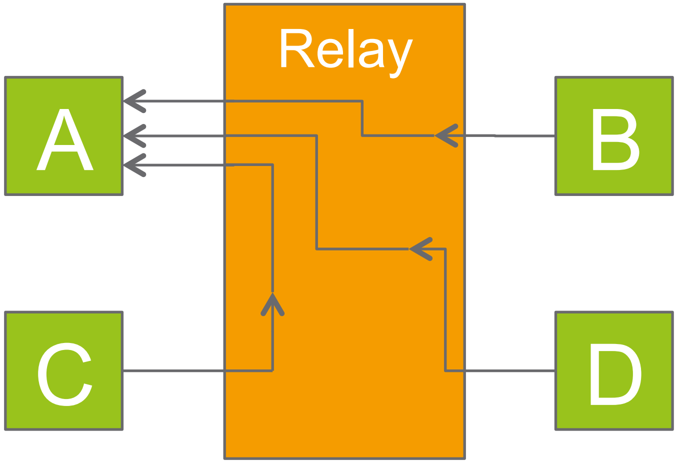

- Description:
  - **One peer transmits only to the Relay, which forwards to multiple peers**
  - Each peer to Relay communication is unicast
  - **Relay creates a conference-wide RTP session**
- Security:
  - **SRTP's regular source authentication can't authenticate peers**
    - **For cryptographic verification TESLA or similar is needed**
  - Confidentiality and Integrity shared with all end-points
  - **Switch need not be trusted with media content**
  - Additional Keying mechanisms likely needed
- Congestion:
  - **Each sender must aggregate receiver statistics reports or requests from all receivers**
  - **All Peers will share available capacity on all paths**
  - Any encoding changes due to congestion will affect all peers
- Identification:
  - Sender's sources map one to one with RTP media streams
- Bandwidth:
  - **Receiver bandwidth will always be the lowest common denominator from all paths**
  - Bandwidth optimizations must occur over whole conference not for individual paths
- Quality:
  - Will be the lowest common denominator based on Capacity and Properties for all Paths
- Complexity:
  - All in the end-points
  - Limitations in end-point directly affect what sender can produce and receiver can accept
  - Conference properties decided by lowest common denominator of peers

Selective Forwarding Switch

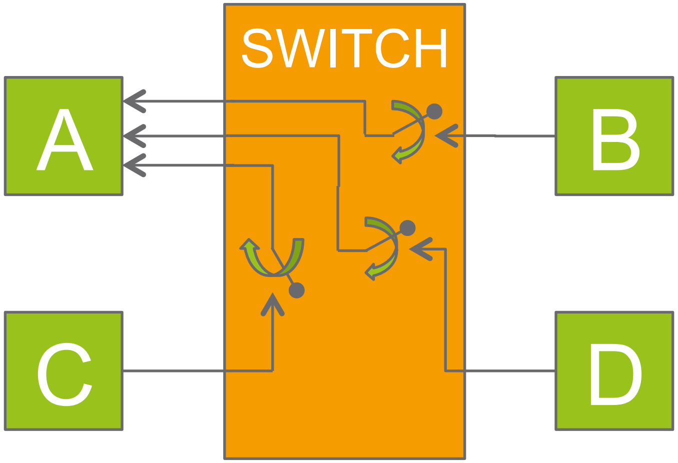

- Description:
  - One peer communicate only with Switch
  - Each peer to Switch communication is unicast
  - **Switch creates a conference-wide RTP session**
  - **Switch turns individual source on and off based on some policy**
  - **Not supported by today's RTP!**
- Security:
  - SRTP's regular source authentication can't authenticate individual peers
    - For cryptographic verification TESLA is needed
  - Confidentiality and Integrity shared with all end-points
  - Switch need not be trusted with media content
  - Additional Keying mechanisms likely needed
  - Switching a source off and later on can break SRTP Roll over Counter
- Congestion:
  - Each sender must aggregate receiver statistics reports or requests from all receivers
  - All Peers will share available capacity on all paths
  - Any encoding changes due to congestion will affect all peers
  - **Reporting and thus congestion detection will be confused by disappearing and reappearing sources**
- Identification:
  - Sender's sources map one to one with RTP media streams
- Bandwidth:
  - Receiver can request media to be tailored to its needs, but will typically also affect other receivers
  - **Which media streams an end-point receives can be individually tailored**
  - Desirable to limit the amount of trade-off based on path capacity
  - Simulcast and scalable encoders can be used to meet different bandwidth needs or requirements
- Quality:
  - Trade-off of end-to-end Path Capacity and Properties between receivers sharing media from the same sender
- Complexity:
  - **Switch has *no* end-point complexity per end-point in the conference, only forwarding logic and tables**

End-point Forwarding

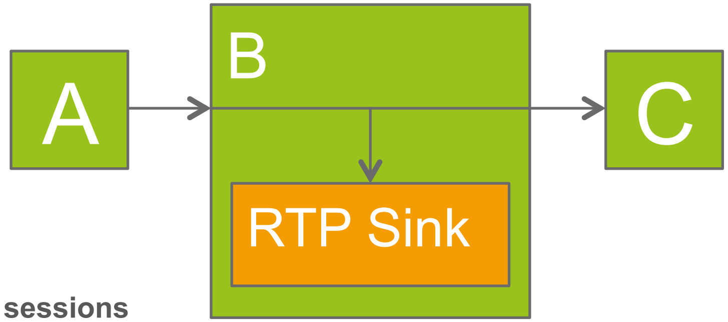

- Description:
  - **One peer communicate only with peer, forwarding to other peers**
  - Each peer to peer communication is unicast
  - **If only forwarding RTP then a common RTP session is created**
  - **If B implements transcoding / RTP mixer functionality multiple RTP sessions can be created (Not further discussed, see Mixers)**
- Security:
  - SRTP's regular source authentication can't authenticate individual peers
    - For cryptographic verification TESLA is needed
  - Confidentiality and Integrity shared with all end-points
  - Additional Keying mechanisms could be used to avoid decryption / encryption cycle in B
- Congestion (from A's perspective):
  - Sender must aggregate receiver statistics reports or requests from all receivers
  - All Peers will share available capacity on shared paths
  - Any encoding changes due to congestion will affect all peers
- Identification:
  - Sender's sources map one to one with RTP media streams
- Bandwidth:
  - Receiver bandwidth will always be the lowest common denominator from all paths
  - Bandwidth optimizations must occur over whole conference not for individual paths
- Quality:
  - Will be the lowest common denominator based on Capacity and Properties for all Paths
- Complexity:
  - All in the end-points, with some added complexity in B
  - Limitations in end-point directly affect what sender can produce and receiver can accept
  - Conference properties decided by lowest common denominator of peers

Any Source Multicast

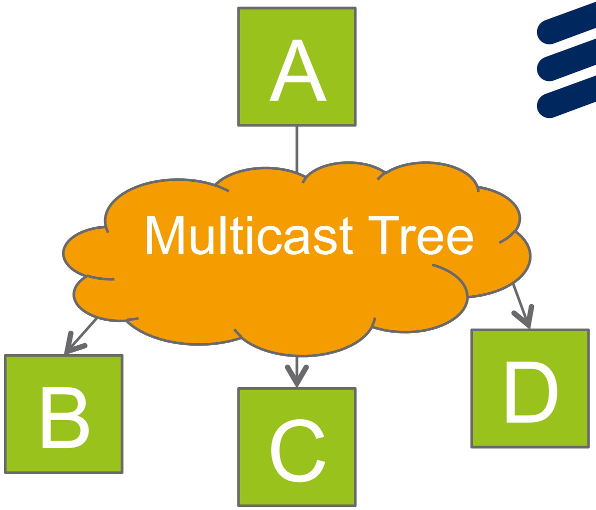

- Description:
  - **One peer communicate with all multicast group members**
  - **Multicast group is a conference-wide RTP session**
- Security:
  - SRTP's regular source authentication can't authenticate individual peers
    - For cryptographic verification TESLA is needed
  - Confidentiality and Integrity shared with all peers
- Congestion:
  - **Each sender must aggregate receiver statistics reports or requests from all receivers**
  - **All Peers will share (single copy) available capacity on all links**
  - Any encoding changes due to congestion will affect all peers
- Identification:
  - Sender's sources map one to one with RTP media streams
- Bandwidth:
  - Receiver bandwidth will always be the lowest common denominator from all paths
  - **Bandwidth optimizations for a single multicast group must occur over whole conference not for individual paths**
  - Bandwidth adaptation can be achieved using multiple multicast groups and simulcast or scalability
- Quality:
  - Will be the lowest common denominator based on Capacity and Properties for all Paths
- Complexity:
  - All in the end-points
  - Limitations in end-point directly affect what sender can produce and receiver can accept
  - Conference properties decided by lowest common denominator of peers

Source Specific Multicast

- Description:
  - A SSM tree enables media delivery to a number of receivers from aggregation point
  - Media sources may be mixed, switched, selected etc. to generate media streams sent over SSM
  - A receiver of the SSM media provides feedback (RTCP) over unicast
  - If a receiver likes to send media it must be sent to media aggregator using separated unicast traffic
- Implications are left as an exercise ;-)

Conclusions

- There are many topologies
  - Most, if not all are valid implementation choices for CLUE systems
- Difficult to select trade-offs to optimize conference
- Do we need to select supported topologies?
- Does CLUE signaling need to take all into consideration?