Some Thoughts On Residential Ethernet Services And Their Definition

Thomas Dineen
Dineen Consulting
Milpitas, CA.
(408) 956-0539 (O)
tdineen@DineenConsulting.com
tdineen@ix.netcom.com
03/02/05

Goals Of This Presentation

- Present an application overview and scope of some proposed ResE work items.

- Introduce a ResE Subscription Protocol (RESP). Discussing the related algorithms and registration protocol.

- Present and discuss related architectural concepts.

- Define ResE related terminology and concepts.

- In regard to work split the content of this presentation would likely be standardized under an IEEE 802.1 PAR.

Application Overview & Scope

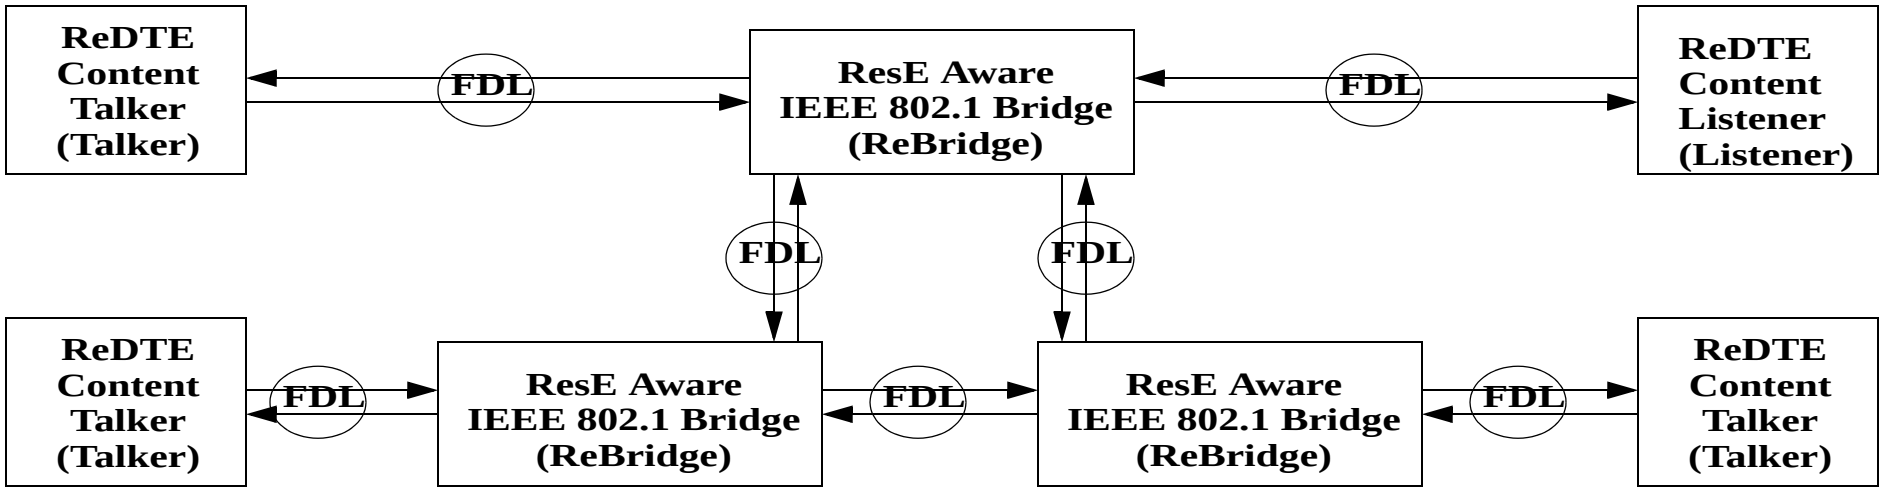

- The proposed architecture supports arbitrary interconnection of ResE Aware End Stations (ReDTEs) and ResE Aware 802.1 Bridges (ReBridges) with up to N-Hops of ReBridges. Where value of N-Hops is application dependant.

- Talkers, Listeners, and ReBridges may be located any where in the Bridged Network or Concatenated Data Link (Catalink).

- FDL - IEEE 802.3 Full Duplex Link or logically equivalent.

- See terminology slides for details.

Application Overview & Scope

- This Architecture assumes the use of 802.3 or logically equivalent Full Duplex links for transit of Content Streams (CS), i.e. No CSMA/CD Half Duplex links are supported.

- Where N-Hops for live entertainment is limited to less than seven hops by application latency requirements, and for recorded entertainment is limited by the seven hop maximum of Spanning Tree.

- 10 MBit/Sec. links are not supported.

- This architecture coexists, without degradation of existing services, within LAN Bridged networks.

- Network Layer Router (IPv4) support is outside of the scope of the proposed architecture.

- This proposed architecture neither explicitly or implicitly requires the use of any Network Layer Standards (IPv4 or IPv6).

- We envision that with the exception of the proposed enhancements the proposed architecture will be fully compatible with IEEE 802, IEEE 802.1, and IEEE 802.3, and will in no way inhibit the interworking of legacy protocols.

Architecture Overview

- Multicast support is specifically limited to one Talker to N Listeners.

- The proposed architecture is comparable with the IEEE 802 and IEEE 802.3 defined address syntax and semantics.

- ResE Content Streams may be point to point or point to multipoint.

- An existing point to point Content Stream may have additional Listeners added to become a point to multipoint Content Stream and visa versa.

- ResE Content Streams shall always use multicast DAs, with bridge forwarding being based on the 802.1 multicast forwarding concept. This facilitates seamless uninterrupted add and drops of ResE Listeners.

- The Multicast DA shall be locally unique, within the LAN scope, or equivalent. This allows the multicast DA to be a unique identifier of the Content Stream within the LAN, and thus allows application of 802.1 multicast filtering. More on this issue later.

- We envision the proposed architecture to be fully compatible with IEEE 802.1 VLANs, Rapid Re-configuration Spanning Tree, and Spanning Tree per VLAN, including finding performance enhancing application of these standards.

Multicast Scope And Pruning

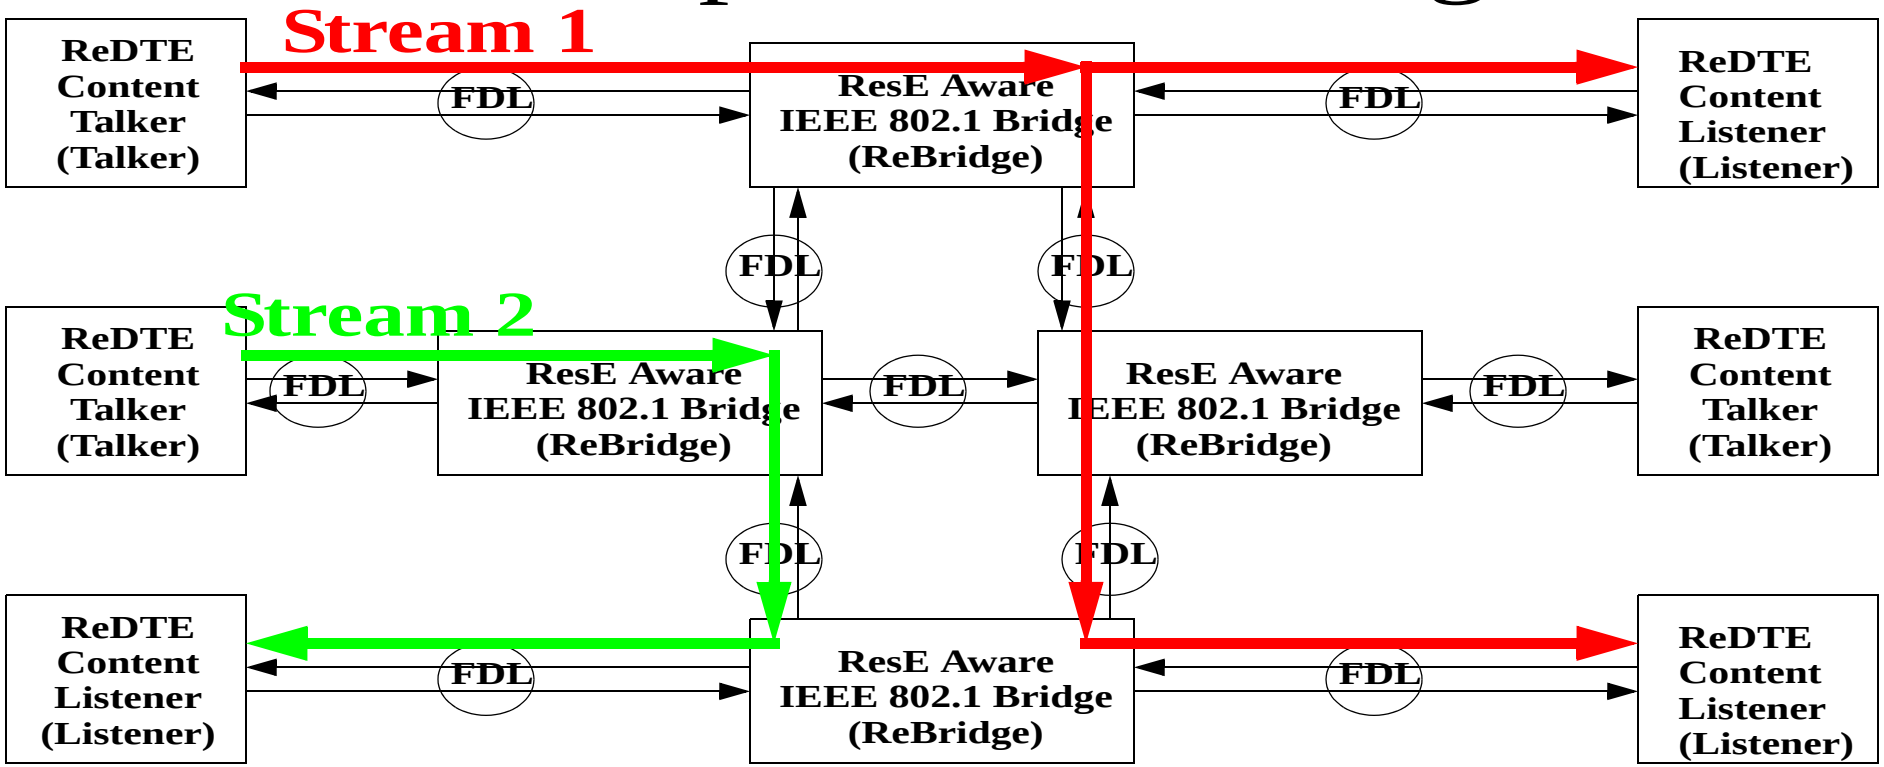

- The 802.1 D,p multicast scope and pruning scheme delivers the ResE Content Streams to just the subscribers thus saving network bandwidth and ReBridge queueing resources.

ResE RESP Overview

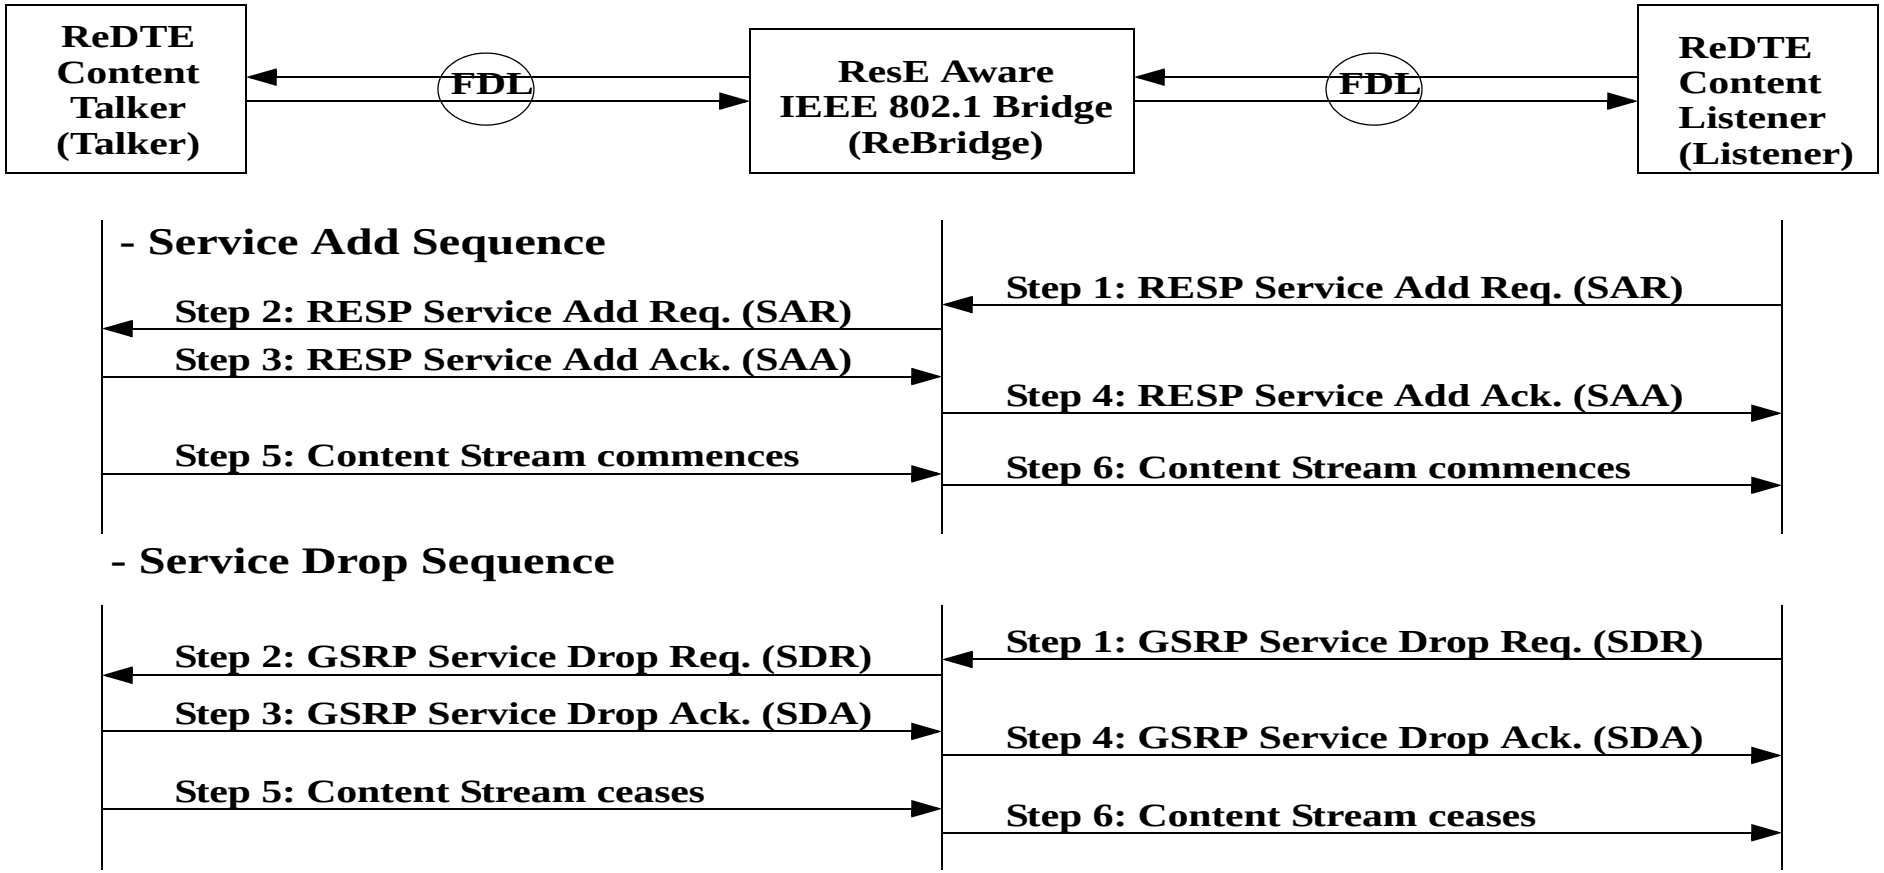

- We propose the definition of a service registration and dergistration protocol, which allows Listeners to subscribe and unsubscribe to ResE Content Channels and their implied ResE Content Streams.

- Where discovery of ResE Content Channels is outside of the scope of this standard.

ResE RESP Overview (1)

- This new protocol will likely be defined as an IEEE 802.1 GARP (MRP?) Application named the “ResE Subscription Protocol (RESP)” and will provide the necessary semantics to register and deregister ResE Content Channels.

- We propose to define a frame based “slow protocol” in support of communication of Add / Drop semantics and their required acknowledgements.

- Some basic RESP Requirements:

- Supports a reliable Service Add / Drop Protocol which allows Listeners to subscribe and unsubscribe on demand to a previously discovered Content Channel.

- Maintaining as required ResE Context State information, in ReBridges and ResE End Stations, in support of the setup, maintenance, and removal of ResE Associations, including information in support of resource reservation requirements.

- Supports Join / Leave Semantics at ReBridges to previously established ResE Associations.

- Define RESP Service Add, Drop, and Acknowledge PDU Formats.

ResE RESP Overview (2)

- Define an internal service interface or an enhancement of an existing internal service interface within ReBridges in support of communication of resource reservation requirements for isochronous queueing. These resource reservation requirements being derived from the ReBridge's stored ResE Context State.

- Support a Service Ping Request and Acknowledge protocol with associated semantics capable of detecting failed or ungracefully removed Re Devices.

- Define a GSRP Service Ping Request (SPR) and Service Ping Acknowledge (SPA) PDU Format, and Context Maintenance Semantics.

- ReBridges send SPR PDUs to connected ReDTEs and ReBridges. On reception the SPR is parsed by the receiver and if valid an SPA response is issued.

- ReBridges will maintain a Watch Dog Timer which is started on transmission of an SPR and stopped on reception of a valid SPA Response.

- If no valid SPA Response is received the Watch Dog Timer expires and the shared ResE Associations then are subject to maintenance action (tear down).

ResE RESP Overview (3)

- Details of the definition of RESP as a GARP (MRP?) based protocol will be presented in future presentations, but first we need to think about the ResE Subscription Service Semantics.

I plan to review the applicability, proposed schedule, and progress of the MRP project and if feasible use MRP as the basis RESP definition.

- From the 801.1ak PAR a bit about MRP:

- 802.1ak Multiple Registration Protocol - This standard specifies protocols, procedures, and managed objects to support the Multiple Attribute Registration Protocol. MRP allows participants in a MRP Application to register attributes with other participants in a Bridged Local Area Network. Two Applications are defined, to register VLANs (MVRP) and Group MAC addresses (MMRP). MVRP will furthermore provide for the rapid healing of network failures without interrupting services to unaffected VLANs.

RESP Service Add Diagram

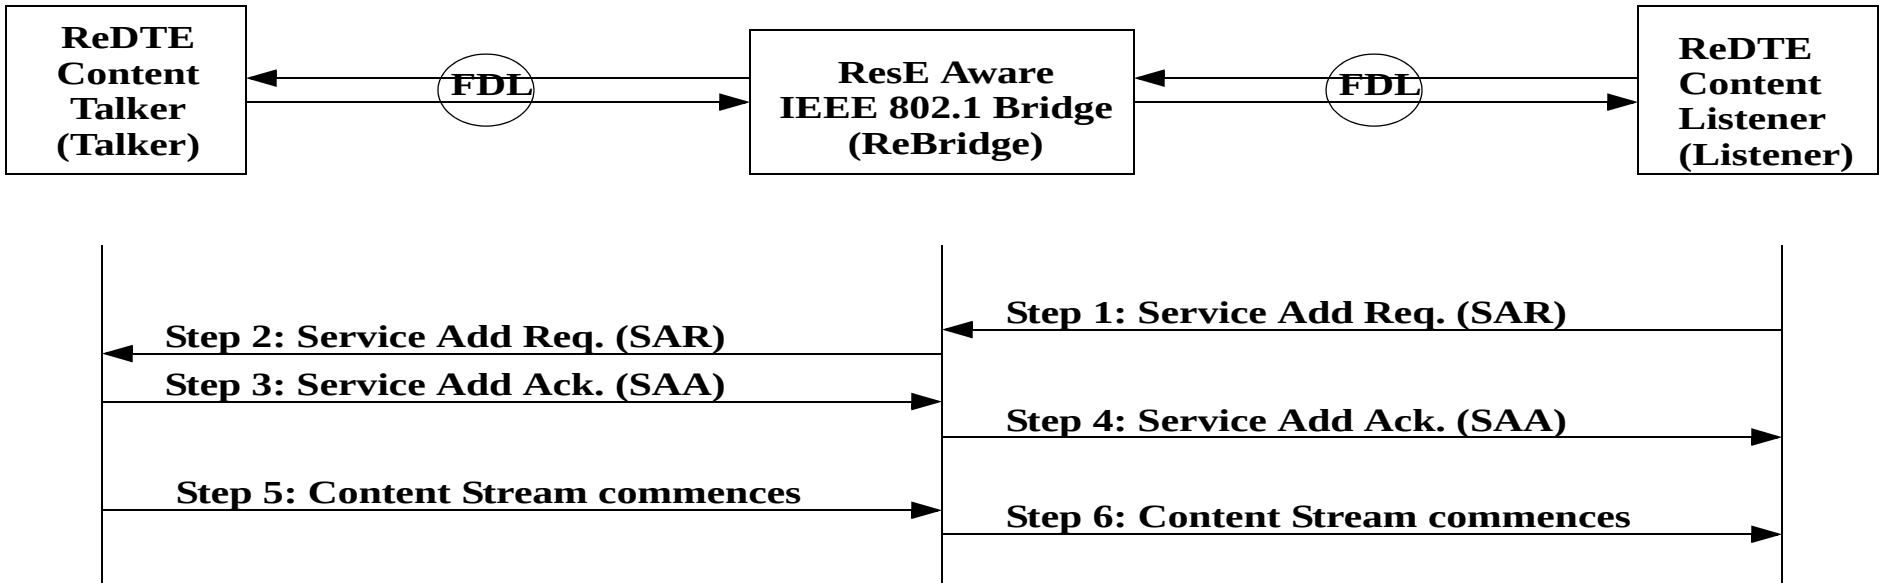

- The GSRP Service Add and Acknowledge exchange allows a ResE Listener to subscribe to a previously discovered ResE Content Channel.

RESP Service Add Description

- Step 1: The ReDTE issues a GSRP Service Add Request (SAR) addressed to the directly attached ReBridge, containing both the desired ResE Content Channel and its selected ListenerPortId. Any required local Listener resources are allocated as pending.

- On reception the ReBridge validates the SAR and confirms that the requested local queueing resources are available. If the required resources are available the ReBridge allocates as pending the requested local resources to the request.

- Step 2: The ReBridge issues a SAR addressed to the requested Content Talker.

- On reception the Content Talker validates the SAR and confirms the requested ResE Service, Channel, and resources are available. If the requested service and resources are available they allocated to the request and the Content Stream begins.

- Step 3: The Content Talker issues a GSRP Service Add Acknowledge (SAA) to the ReBridge.

- Step 4: On SAA Reception the ReBridge validates the SAA. If valid it moves requested pending resources to allocated and issues an SAA to the Listener.

RESP Service Add Description (1)

- On SAA reception the Listener validates the SAA. If valid it moves the requested pending resources to allocated and reception of the Content Stream (CS) begins.

- If any device in the above sequence lacks the requested resources it simply responds with a RESP Service Add NAck (SAN), which is propagated back to the requesting Listener, releasing pending resources.

- Error handling is in the purview of higher layer protocols.

- Listeners are required to Add when the customer requests a channel. Note: If a Customer changes the requested channel on a Listener the device first Drops the current subscribed channel and then performs the Add of the new selected channel.

GSRP Service Drop Diagram

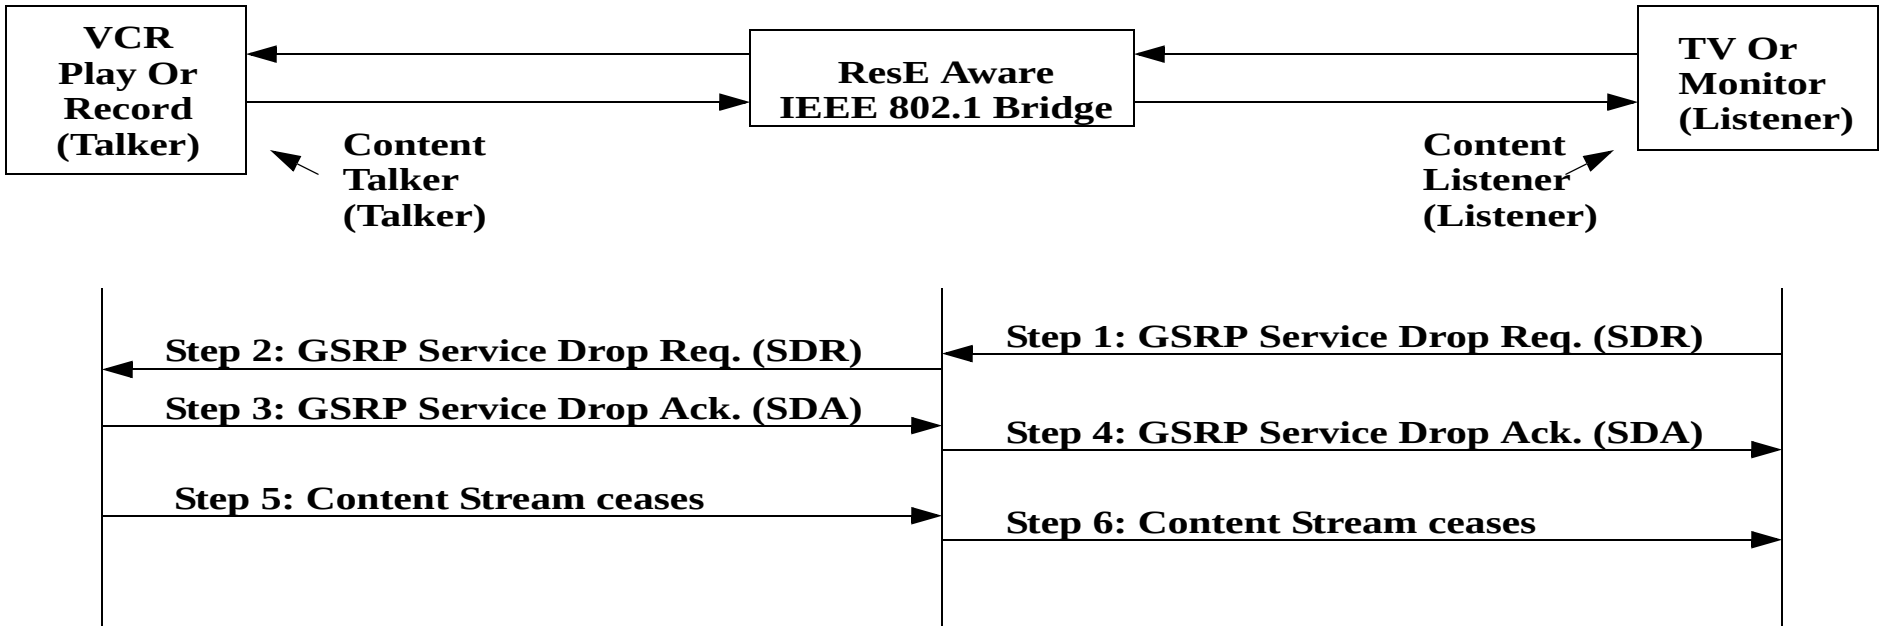

- The GSRP Service Drop and Acknowledge exchange allows a Rese Listener to unsubscribe to a previously discovered and subscribed ResE Content Channel.

RESP Service Drop Description

- Step 1: The Listener issues a GSRP Service Drop Request (SDR) addressed to its directly attached ReBridge, containing the TalkerPortId and ListenerPortId of the previously established Re Association. Any Listener resources which are previously allocated to the CS are moved to pending.

- Step 2: On SDR reception the ReBridge validates the SDR. If valid any resources which are previously allocated to the CS are moved to pending. The ReBridge issues a SDR addressed to the Talker.

- Step 3: On SDR reception the Talker validates the SDR. If valid the CS is terminated, and previously allocated resources are moved to deallocated. The Talker issues a GSRP Service Drop Acknowledge (SDA) to the ReBridge.

- Step 4: On SDA reception the ReBridge validates the SDA. If valid it moves pending resources to deallocated, and issues an SDA to the Listener.

- Step 5: On SDA reception the Listener validates the SDA. If valid it moves pending resources to deallocated, and the ResE Channel with its associated ResE Content Stream is considered gracefully dropped.

RESP Service Drop Description (1)

- When dropping Content Streams (CS) Talkers and ReBridges must consider that a CS may have multiple subscribers consequently an acknowledged SDR command may cause in effect a pruning of a single Listener from a multicast distribution tree which is composed of a single Talker, one or more ReBridges, and several Listeners.

- A Listener issuing an SDR may fail to receive an acknowledge due to device failure or ungraceful power down. As mentioned previously we will need to include a keep alive or heart beat protocol based on the SPR and SPA concept. On Watch Dog Time-out stale resources will become deallocated and thus available for reuse.

Locally Unique Multicast

- There have been three choices proposed so far as to the locally unique multicast addressing problem:

1) Lower 24 bits (OUI) (Octets 0 through 2 per Std. 802) of Multicast DA are as defined by IEEE 802 are unchanged. Upper 24 bits are a locally generated pseudo random value. There is a 1 in 2 to the 24 chance of a collision. This just is fine. A collision in this context just means that two multicast Streams are sharing a multicast address identifier thus the PDUs of both Streams are delivered to both sets of subscribing ResE Listeners. So all ResE Listeners just filter on the TalkerPortIds they have subscribed to. The driver would likely be required to perform this function. BTW NICs already do this for multicasts they have subscribed to! This alternative would require either a single reserved ReOUI or the use of locally administered multicast addresses, to prevent collision with legacy address assignments.

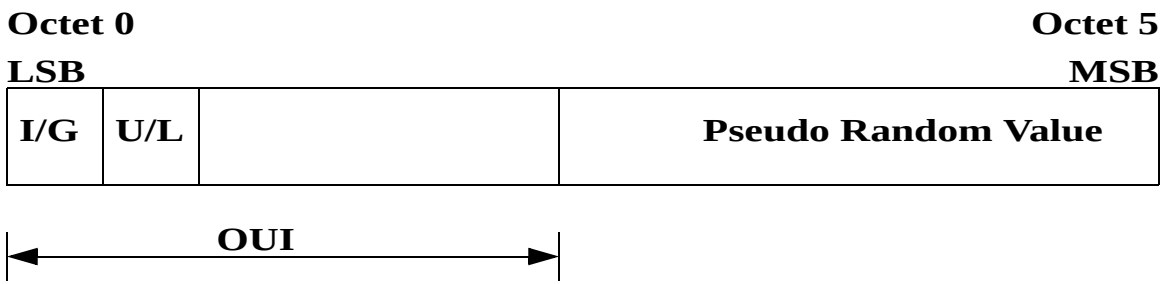

Locally Unique Multicast (1)

2) Use a central Multicast Address Allocation Server which is probably located on an ReBridge, and probably is the Re Timing Root ReBridge. Note: This server just serves up locally unique multicast addresses on request. It not like an ATM BUS. This alternative would require either a single reserved ReOUI or the use of locally administered multicast addresses, to prevent collision with legacy address assignments.

3) Modify the ReBridge forwarding protocol to forward ResE Content Stream PDUs based on more than the current multicast DA. To form a unique a bridge forwarding identifier concatenate the Talker's Source Address and the TalkerPort. This forwarding identifier replaces the multicast DA when indexing the bridge forwarding table. This alternative would require either a single reserved ReOUI or the use of locally administered multicast addresses, to prevent collision with legacy address assignments.

ForwardingIdentifier := {TalkerTalkerAdr,TalkerPort};

Terminology

- RESP - ResE Subscription Protocol.

- ReDTE - A ResE aware computer end station or server.

- ReBridge - A ResE Aware 802.1 Bridges.

- FDL - IEEE 802.3 Full Duplex Link or logically equivalent.

- N-Hops - The number of ReBridge Hops in a Content Stream Path (CSP).

- ReC - ResE Entertainment Content any live or prerecorded entertainment product which can be represented in digital form and transferred as packetized information.

- ResE Content Talker (Talker) - A Talker or source of ResE Entertainment Content such as a VCR playing, a cable set top box, or a musical instrument.

- ResE Content Listener (Listener) - A Listener or sink of ResE Entertainment Content such as a VCR Recording, a TV, or speakers.

- ResE Content Stream (CS) - Any Layer 2 flow of ResE Entertainment Content from any higher layer application formatted as a series of 802.3 frames traveling from a Talker to a Listener, via a ResE Content Path. RESP PDUs are not part of CSs.

Terminology Continued (1)

- ReCP - ResE Content Path - A logical concatenation of point to point links and ResE Aware Bridges defining a continuous path from a ResE Talker to a ResE Listener.

- ReAud - ResE Audience - The set of ResE Listeners subscribing to a ResE Content Stream and sharing a ResE Association.

ROADS - Running Out Of Addresses Problem (IETF Terminology).

- Catalink - Concatenated Data Link - Bridges transparently concatenate two or more independent Data Links to form a Concatenated Data Link.

- RePort - ResE Port - Any IEEE 802.3 Port operating in Full Duplex mode which contains an compliant instance of the logic specified by the proposed Re Standards.

- RePortId - ResE Port Identifier - a 16 bit port association identifier, used within the context of Talker and Listener Port Identifiers defined below.

- ReCC - ResE Content Channel - A 64 bit identifier which Identifies a unique instance of entertainment content available to subscribers via an Re LAN Network. ReCC identifiers are unique within a ResE Network. Discovery of ResE Content channels is outside the scope of this standard.

Terminology Continued (2)

- TalkerPortId - Talker ResE Port Identifier - The Talker Port Identifier is unique to the talker port and assigned by the Talker Port or its surrogates when a ResE Association is established. Within Talkers the TalkerPortId provides an unique index into the Talkers's ResE Association Table.

- ListenerPortId - Listener ResE Port Identifier - The Listener Port Identifier is unique to the Listener Port and assigned by the Listener Port or its surrogates when a ResE Association is established. Within Listeners the ListenerPortId provides an unique index into the Listener's ResE Association Table.

- ReAS - ResE Association - A logical association including ResE Context State information between a ResE Talker, one or more ResE Listeners, and one or more intervening ReBridges, comprising a ResE Content Stream.

- ReCON - ResE Context - The ResE State information stored within a ReDTE or ReBridge which represents its portion of the ResE Association State for an active ResE Association.