Layer 2 “RSVP”: MSRP? SBM? Other?

Choosing the right method to reserve bandwidth and measure latency across bridged LANs

Rev. 1

Norman Finn, Cisco Systems

Note

- This presentation can be found at: http://www.ieee802.org/1/files/public/docs2007/at-nfinn-msrp-sbm-rsvp-0712-v1.pdf.

The Audio Video Bridging problem statement

Home network environment

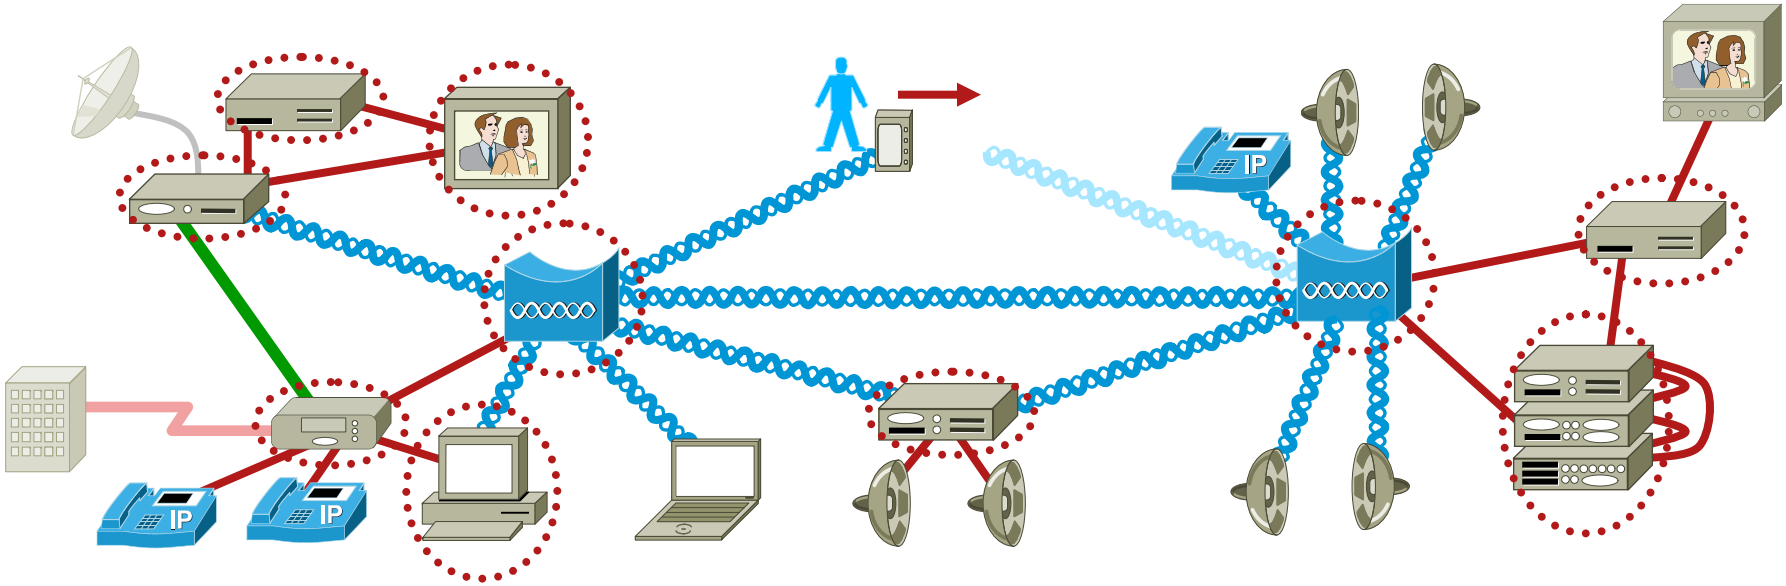

- In a home or small studio, there may be many Ethernet-like links: 802.3, 802.11, MoCA, Ether/DSL, etc.
- One way to achieve connectivity within the home is that every device with multiple links is an 802.1 bridge.
- This does not preclude routing (e.g., over the DSL link.)

Arbitrary connectivity

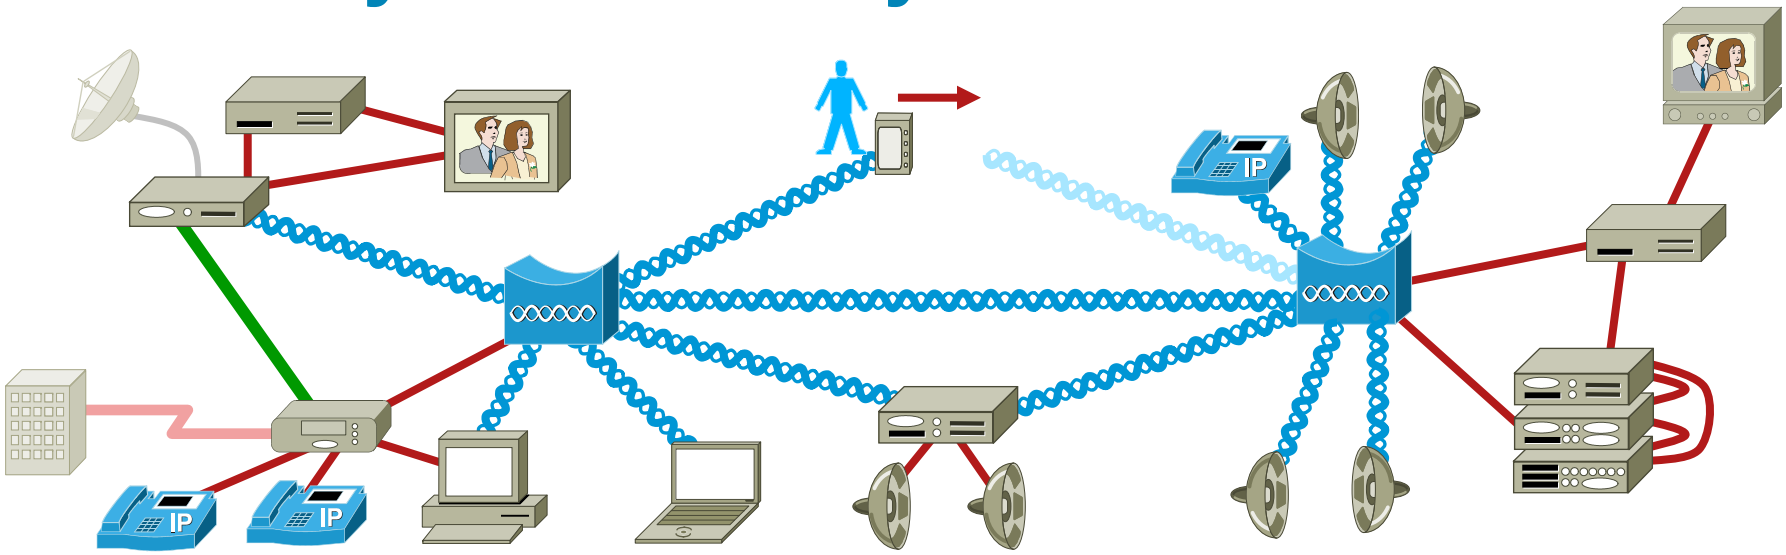

- Note that no structure is imposed on the wiring plan.
- This home network example shows islands of wired connectivity around a wireless backbone. (A single AP would be more common, of course.)
- A studio example would be more heavily wired, because it has much higher bandwidth requirements.

No masters, no slaves

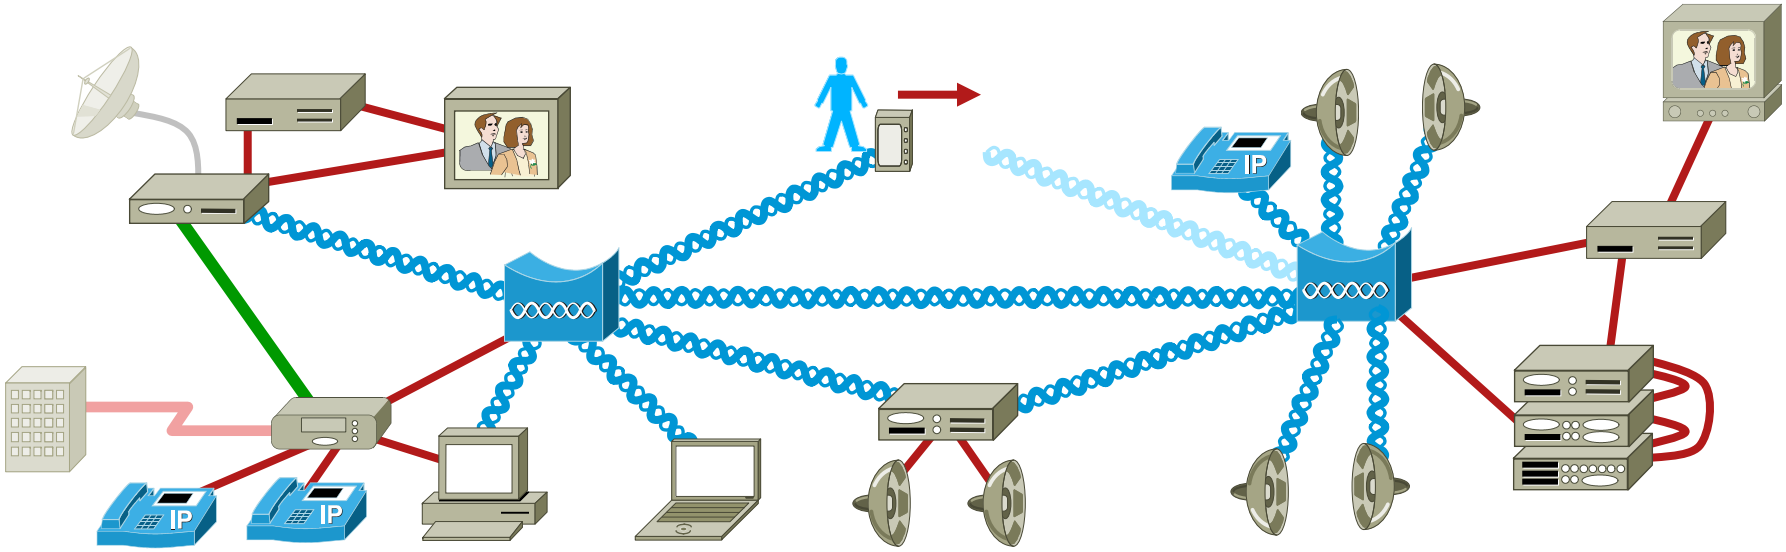

- IEEE 802.1 assumes that no godbox supervises the network. (Perhaps a religious issue, but one with a sound technical basis.)
- The network must be 100% plug-and-play.

AVB Task Group's task list

- AV devices require synchronization of their clocks to within 10 µs. (Projects 802.1AS and 1588 v2.)
- A stream reservation protocol is required so that stations can reserve bandwidth and measure latency across the bridged network. (Project 802.1Qat.)
- A better output queue selection algorithm than strict priority selection is required to support stream reservation. (Project 802.1Qav.)
- IEEE 802.1 and IEEE 802.11 standards require revision to define a device that is both a bridge and a wireless station. (In discussion between 802.1 and 802.11.)
- A profile for bridges and stations, requiring the above protocols, better-than-minimum conformance to 802.1Q, and some altered default configuration parameters, is required to ensure plug-and-play operation and support of the target applications. (Project 802.1BA.)

Serial vs. Parallel operation

Subnet Bandwidth Manager: Serial

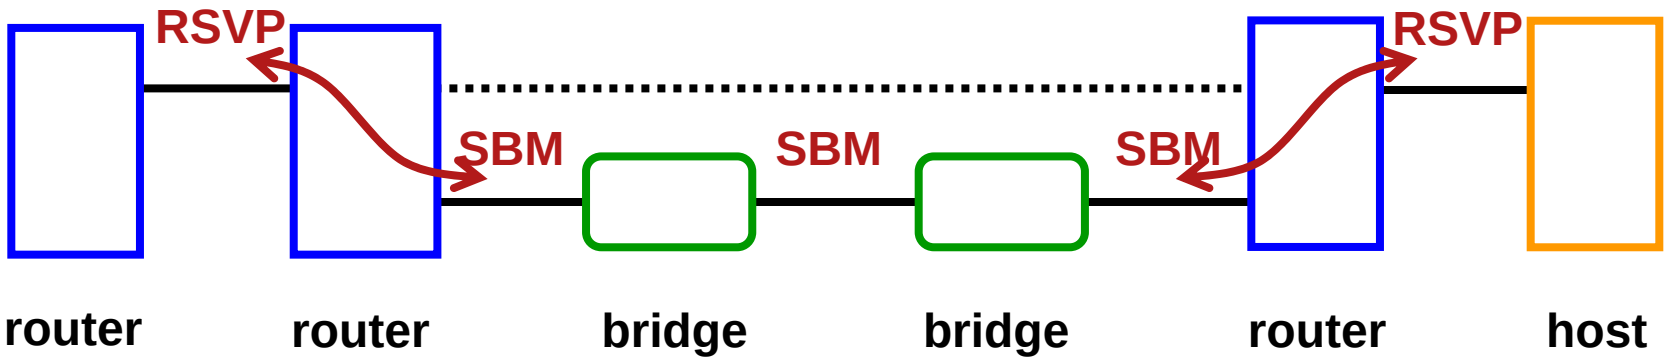

As this author (a bridge person) understands:

- RSVP messages pass from router to router until they reach an SBM-capable bridge.
- Layer 2 information is added to the RSVP message, transforming it into an SBM message.
- At the end of the bridged network, the SBM message is transformed back to an RSVP message and continues on its way.

Preferred for layer independence: Parallel

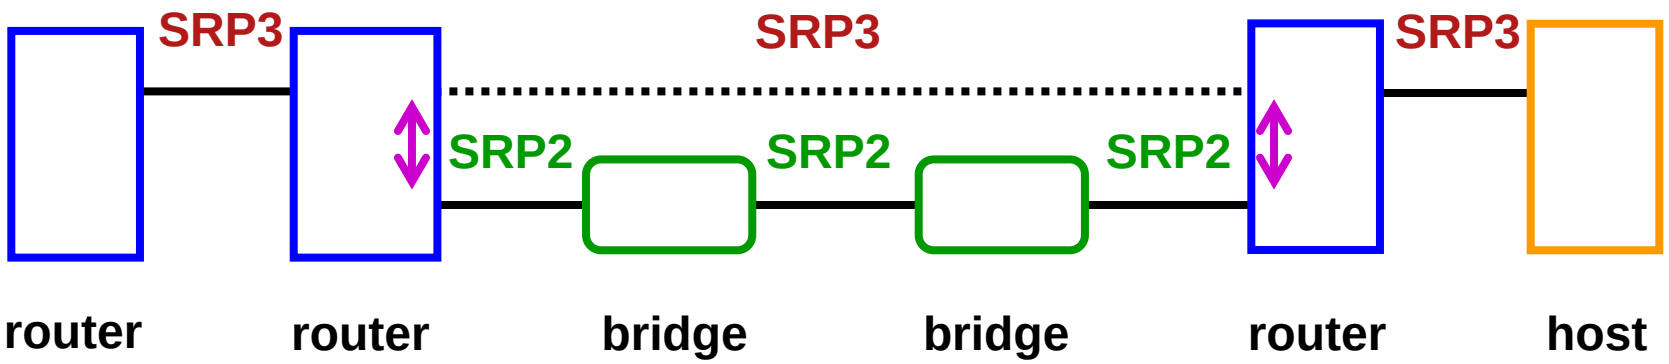

- Bridges and higher layers each have their own stream reservation protocols (SRP2 and SRP3, above).
- A router or host that handles both protocols must resolve the permission and/or rejection of reservations by the two protocols.
- At least one of SRP2 or SRP3 has to carry an ID tag allowing SRP2 and SRP3 reservations to be associated.

P802.1Qat direction, at the moment

P802.1Qat MSRP snapshot

As of this writing, P802.1Qat (Multiple Stream Registration Protocol, MSRP) is headed in the following direction:

- IEEE 802.1 bridges operate independently of any higher layer protocols, with a minimum of layer interconnects (e.g., ARP).
- Stream reservation must not be bound to any higher layer protocols, nor require any higher layer protocols to function.
- MSRP must provide a facility that is useful to and compatible with known higher layer stream reservation protocols.
- RSVP is the only higher layer protocol considered.

P802.1Qat MSRP snapshot

- MSRP does not assume that the paths from the talkers to the listeners have been established (e.g., by IGMP, PIM, GMRP, or MMRP) before the stream reservation is made; reservation and L2 path establishment is done simultaneously by MSRP.
- A subset of the model for data flow used by RSVP will be used by MSRP. Namely:

  A talker offers a stream (Path message information). This is propagated through the network by the bridges.

  Listeners respond (Reservation message information). Responses result in building the tree of properly configured queues from the listeners towards the talker.

  Both talker and listeners are notified of failures.

P802.1Qat MSRP snapshot

- MSRP collects worst-case latency information along the talker-to-listener direction. It is TBD whether any latency information is returned to the talker.
- There is a perceived requirement that some details be provided when a reservation fails.
- MSRP uses both max bytes per second and max frames per second to characterize a flow.
- Multicast MAC address collisions (two flows with same L2 multicast destination address) are an issue; a vendor may have to choose between building a bridge that wastes bandwidth, and a bridge that peeks at layer 3.
- MSRP will handle the multiple-talkers-but-one-at-a-time reservation problem. This is useful in the studio when switching among sources for a single stream.

P802.1Qat MSRP snapshot

- MSRP is not intending to use the RSVP method for transmitting data, but will be based on the MRP model.
- Bridges already implement the Multiple Registration Protocol (MRP) for registering directional information, including MAC addresses and VLAN IDs.

  MMRP prunes the propagation of frames to registered destination MAC addresses.

  MVRP prunes the propagation of frames to registered VLANs.
- Useful MRP features:

  As many declarations per transmitted frame as fit.

  Per-port per-application timers, instead of per-registration timers.

  Provision for anti-chatty and hold-down during netquake.

  Information (though not MRP PDUs) follow the active network topology.

AVB questions for RSVP/SBM experts

Questions

- Are we wrong in rejecting SBM because of its relationship to RSVP?
- What features of RSVP should we take special note of?
- What might we do better than RSVP?
- Who defines RSVP/MSRP interworking? Where?