



Wireless Bridges


Problems with Integrating bridging into 802.11 access points and stations

Norman Finn

Note

- This presentation can be found at:
<http://www.ieee802.org/1/files/public/docs2007/avb-nfinn-wireless-bridges-071307.pdf>.
- This presentation is related to an earlier contribution, “Point-to-Multipoint Bridging”,
<http://www.ieee802.org/1/files/public/docs2007/avb-nfinn-point-to-multipoint-bridging-061307.pdf>.

Executive summary



- In a home or small studio, there may be many Ethernet-like links: 802.3, 802.11, MoCA, Ether/DSL, etc.
- To ensure connectivity, every device with multiple links needs to be an 802.1 bridge.
- The IEEE 802 standards do not support an 802.11 wireless station that is also a bridge.

Object

- The object of this document is to illustrate the gaps in the IEEE 802 standards with regard to 802.11 access points and stations, 802.3 media, and 802.11 bridges.
- These gaps, if not corrected, will make it impossible for the Audio Video Bridging Task Group to achieve its goals.
- This document strives to hold to the current architectures of both 802.1 and 802.11, so that implementing bridging in 802.11 devices need not constrain either Working Group's ability to continue to grow its base of standards.

Assumptions

- 802.11 ad hoc mode is not to be used; all stations and access points operate in **infrastructure** mode.
- **Bridges are not confined to wired connections**, or even to access points; a bridge could be, for example, a personal computer with any number of wired ports, 802.1b station interfaces, and/or 802.1g station interfaces. A network can have any number of interior wireless LANs.
- **Every device** in the bridged network with **more than one port** is a **bridge**. (Routers are beyond the scope of this document.)

The problems include

- **Interruptions in reserved priority flows** that inevitably occur due to unnecessary flooding of data streams, in the absence of solutions to the following problems:

The inability of a station bridge to **suppress reflected multicasts**, yet **learn** MAC addresses from bridged multicasts;


The inability of a station bridge to **suppress reflected unknown unicasts**, yet **learn** MAC addresses from bridged unicasts; and

The **lack of a definition** of how the Spanning Tree and other layer 2 control protocols work in an **access point bridge**, in the presence of station bridges and wireless links to other access point bridges.




A problematical relocation scenario

Problematical relocation scenario




- TV has a single Ethernet port.
- The TV is **playing a unicast program** from the DVD.

Problematical relocation scenario




- The TV is moved to another wired connection.
- The DVD does not, according to the standards, forget that the TV is on the now-disconnected port; it **blackholes** any traffic for the TV. (This is purposeful – it **prevents needless flooding** in the usual case that the TV is no longer reachable through the network.)

Problematical relocation scenario




- If bridge B sends a **Topology Change Notification (TCN)** when the TV associates with it, then all bridges **unlearn** most of their addresses.

Problematical relocation scenario



- And fairly quickly, they all **relearn** the location of all of the devices.
- The TV **program resumes** after a short interruption.

Problematical relocation scenario



- But, until the addresses are learned, **all traffic is flooded everywhere**; not just the DVD–TV traffic, *but all traffic*.
- This flooding will, most likely, result in **glitches in any or all of the audio and video flows** in the network.

Problematical relocation scenario


- This problem is **not new**; it exists in wired networks.
- However, this problem has a **trivial solution**, one that is widely implemented:

When a bridge makes a connection to an “access port”, that is, a port that is configured to be unlikely to be connected to another bridge, it does **not generate a TCN**, and thus, does not interrupt the network with an unnecessary flood of unknown unicasts.

When a station makes a new wired connection, it **sends** a number of multicast and/or **broadcast frames**, e.g., ARP requests.

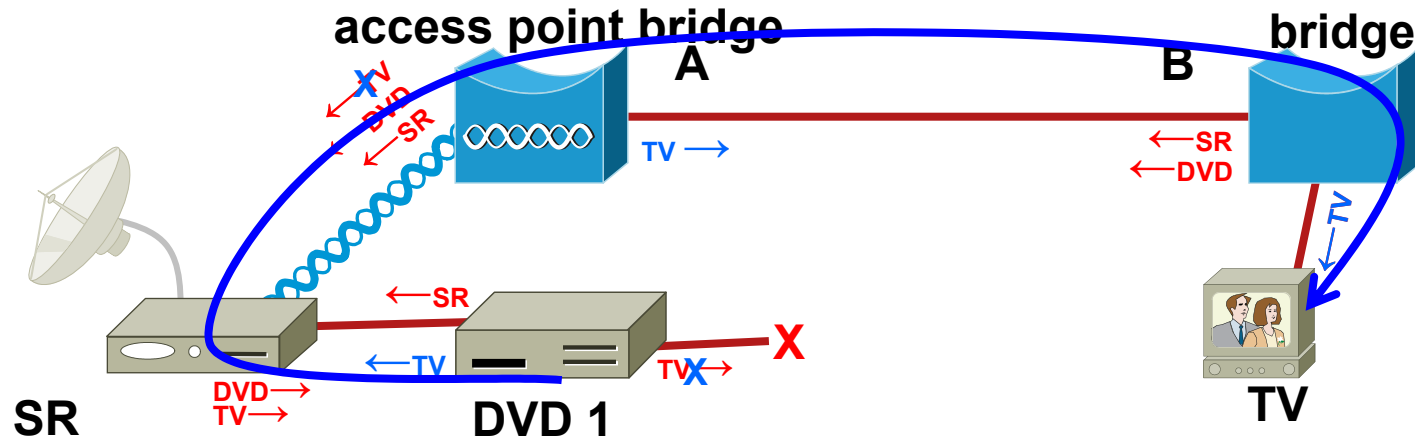
These multicasts/broadcasts are seen by all of the bridges in the network, so they **immediately relearn** the station’s new location.

Problematical relocation scenario




- Applying the current solution, bridge B does not send a TCN, and instead, the TV sends **some broadcasts** and/or multicasts.

Problematical relocation scenario



- What **should** happen: This causes all ports to **relearn** the TV's position.
- The **TV program** resumes.
- There is no flooding burst; the DVD goes from blackholing to transmitting on the correct port.

Problematical relocation scenario




- What **does** happen: Station bridge SR “**knows**” where the TV is, so suppresses its broadcasts.
- So, the TV program does not resume until the SR and TV **time out their MAC table entries** (default is 5 minutes), and start flooding.
- Then, there is a short **flooding burst**, that can **disrupt all flows in the network**, until the TV can respond, and the SR and DVD learn its new location.




Broadcast reflection suppression

Stations and multicasts on 802.11




- If station C **sends a broadcast**, the access point **reflects** the frame back to the wireless medium, and all bridges relay that frame.
- All stations receive the **reflected broadcast**.
- Station C **discards** the reflected broadcast, on the grounds that the **frame's source address is C's address**.

Bridges and multicasts on a Fat Yellow Coax




- Z sends a broadcast. D relays it to the shared medium.
- A, B, and C receive it, A relays it to X.
- Bridges A, B, and D learn the direction to source Z.
- D does not see the frame reflected back towards itself.

Bridges and multicasts on a Fat Yellow Coax



- Z **moves** from D to B, and **sends another broadcast**.
- A, B, and D receive it, A relays it to X.
- Bridges B and D **learn the new direction** to source Z.

Bridges and multicasts on 802.11



- Z sends a broadcast. D relays it to the access point.
- Access Point reflects broadcast to B, C, and D, and relays it to X.
- Bridges A, B, and D learn the direction to source Z.
- D does see the reflected frame and must suppress it, but on what grounds does it suppress the reflection?

Bridges and multicasts on 802.11

- In the current 802.11, the broadcast frame sent by D has three addresses:

Receiver Address: A, the access point.

Transmitter/Source Address: Z, the originator

Destination Address: Broadcast

- The reflected frame also has three addresses:

Receiver/Destination Address: Broadcast

Transmitter Address: A, the access point

Source Address: Z, the originator

Bridges and multicasts on 802.11

- A station, (e.g. C, in this example) can discard its own reflected broadcasts/multicasts by examining the Source Address of the reflected frame, and discarding frames that have its own source.
- But, Bridge D cannot do this, because it may not know for sure whether Z is attached to it, or whether Z has moved behind Bridge B, and must be re-learned, as in the example of the Fat Yellow Coax.

Z may not be logged in with 802.1X.

Z may be behind a chain of bridges.

Z may be behind yet another access point.


Bridges and multicasts on 802.11

- In some network topologies Bridge D can suppress the reflected broadcast on the grounds that it “knows” that Z is behind it.
- But, in a bridged environment, such knowledge is unreliable, because the network topology can change.
- Identifying reflected frames based on the frames recently sent by the station bridge is problematical, because reflected frames can be held in priority queues in the access point before being reflected.



Unicast reflection suppression

Bridges and unicasts on 802.11



- Suppose Station Z sends a **unicast** to Station Y, but no one knows where Y is? Bridge D **floods the frame everywhere**, including towards the access point.
- Access point bridge A must reflect the unicast frame down, since it doesn't know where Y is, either.
- How does D know to **not flood** the reflected unicast?

Bridges and unicasts on 802.11

- Bridge D already flooded the frame to all of its local ports.
- The access point can always transmit unknown unicasts down towards all station bridges for flooding.
- If Bridge D flooded the reflected frame, and if Y were actually attached to D, then Y would get two copies of the frame, which is unusual behavior for a bridged network.

Bridges and unicasts on 802.11

- But, Bridge D does not necessarily know whether a given destination address that is or is not in its own Filtering Database is or is not in the access point bridge's FDB, so it does not necessarily know whether any given frame it sends will or will not be reflected back to it.
- The only way that station bridge D knows, within the current standards, to not flood the reflected frame, is the same as for multicasts – either:

The source address must be “known” to the bridge, which is **unreliable**; or

The frame must be compared to a list of recently-sent frames, which, because of queuing in the access point, is **impractical**.



Separated Mode or Combined Mode?

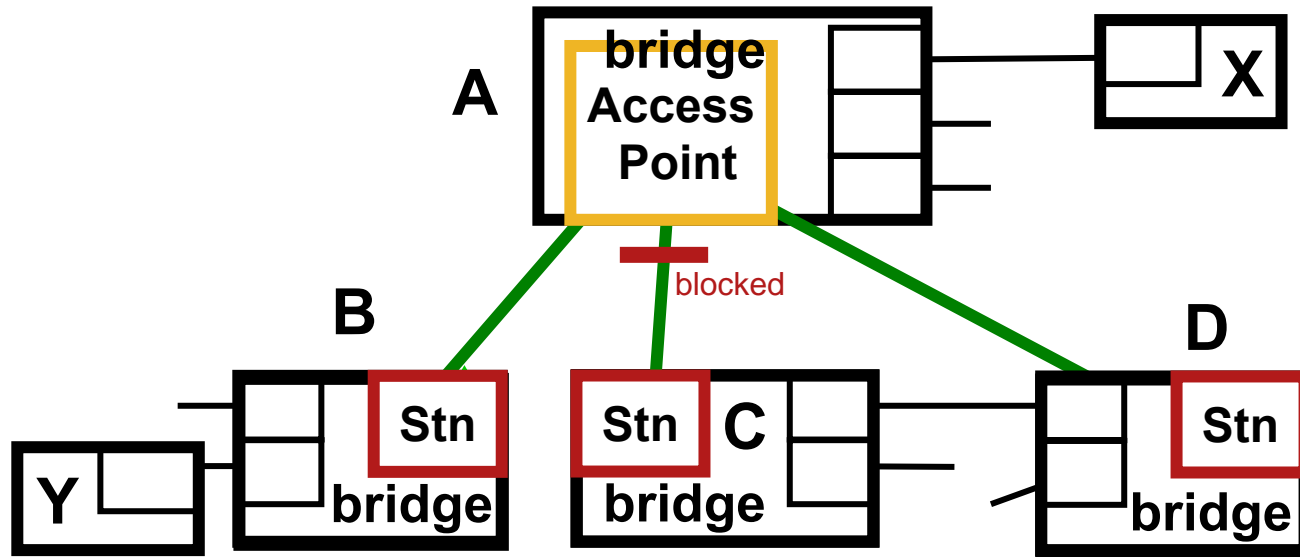
Separated Mode or Combined Mode

- “Point-to-Multipoint Bridging”, <http://www.ieee802.org/1/files/public/docs2007/avb-nfinn-point-to-multipoint-bridging-061307.pdf>, defines two models for how spanning tree can use point-to-multipoint links, such as the wireless medium by which an access point transmits data to its stations:

Separated Mode: Each point-to-point link (association with an individual station) is treated as a separate link from the standpoint of the Spanning Tree Protocols; each can be blocked.

Combined Mode: An access point and all of its associated stations behave, to the Spanning Tree Protocol, as if they are a single shared medium; either all stations are blocked from the access point, or none are.

Separated Mode



- For an access point operating in Separated Mode, spanning tree can block the link individual stations.
- When access point bridge A sends a multicast or floods an unknown unicast, there must be a means to **prevent station bridge C from accepting and relaying it.**

Separated Mode

- **Therefore, we must either:**

- Provide a means to tell each station bridge whether the access point's end of the link is blocked, so it can discard the blocked frames it receives;

- Provide a means to distribute multicasts and flooded unicasts to a subset of the station bridges; or

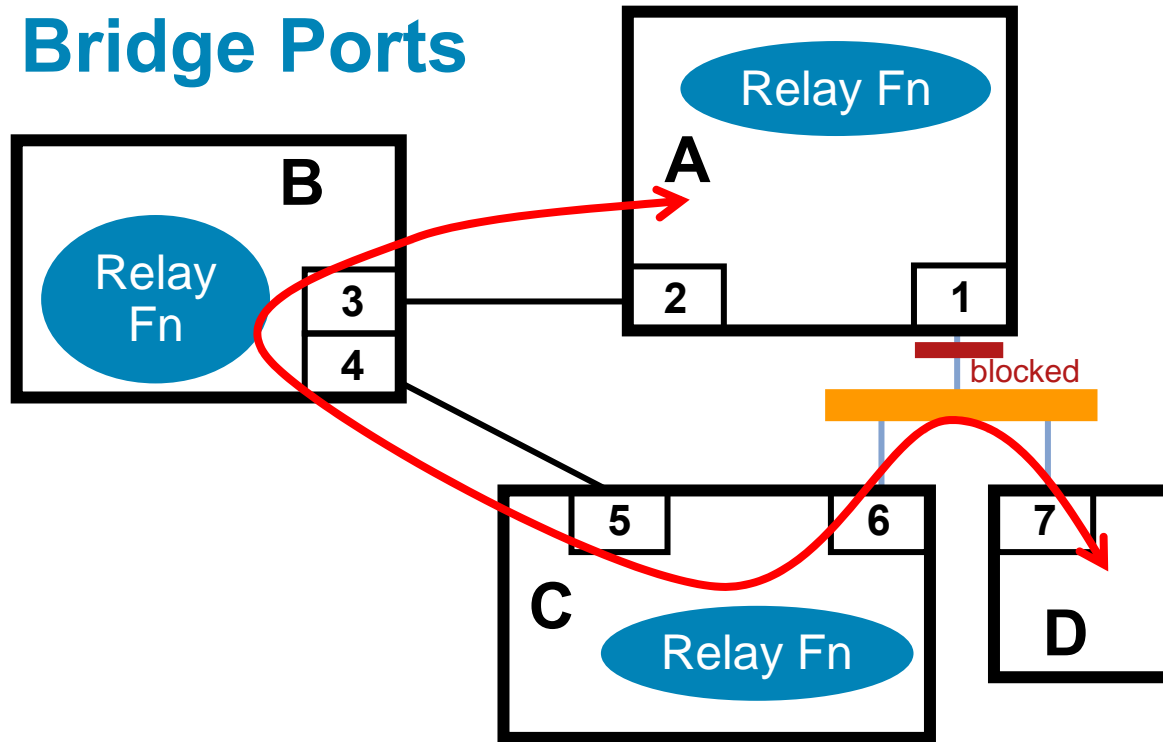
- Replicate and distribute all multicasts and flooded unicasts to one station at a time.

- Any solution is also constrained by the fact that, if the Multiple Spanning Tree Protocol is used, different VLANs can be blocked on different sets of ports.

Combined Mode

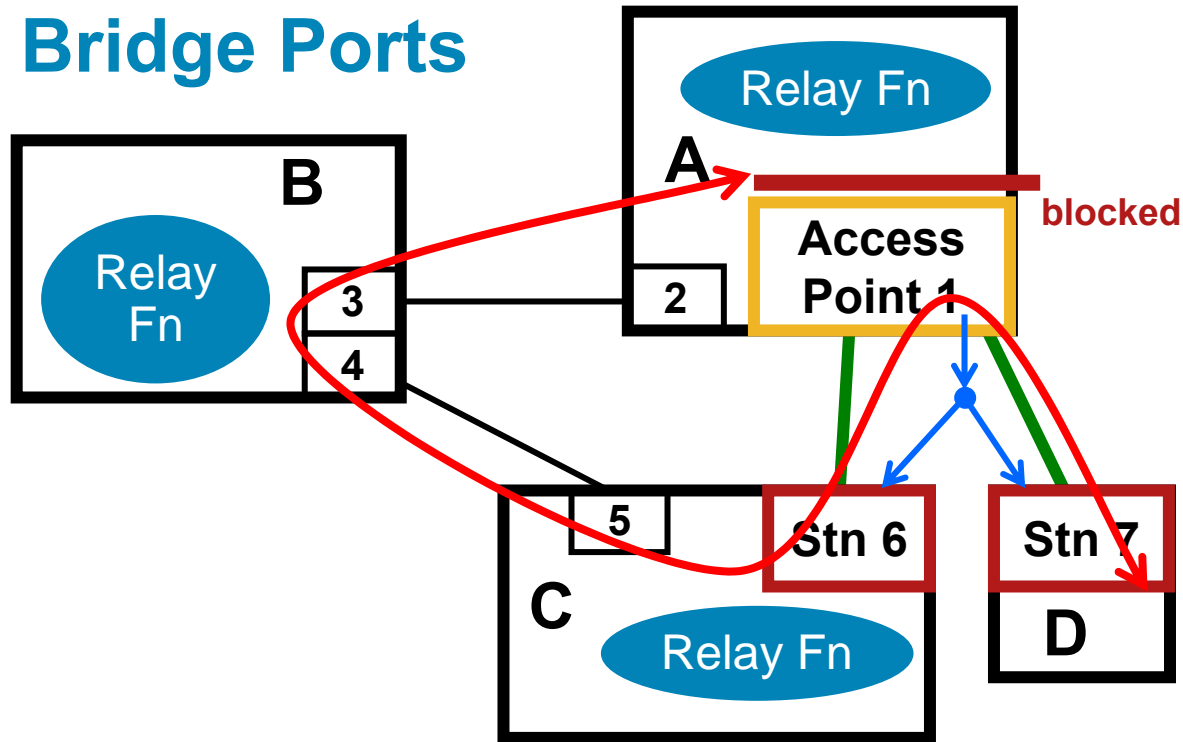
- In **Combined Mode**, the wireless medium is treated as a single emulated LAN; anything sent to or from one station is, in theory, visible to all other stations.

Combined Mode Blocked Bridge Ports



- A shared medium connects Bridge Ports 1, 6, and 7.
- The loop must be broken by blocking a Bridge Port.
- Let's say 1 is blocked. The **arrow** shows the A—D path.

Combined Mode Blocked Bridge Ports



- Bridge **A** in Combined Mode. Perhaps because the 3—2 link is slow, the access point is blocked.
- Access Point **A** must relay data from **C** to **D**, even though its port is blocked.

Combined Mode Blocked Bridge Ports

- Blocking a station bridge's wireless port is easy; the bridge stops data between its wireless medium and its wired connections.
- Blocking the wireless side of an access point bridge is somewhat more complex, since the **access point still forwards data** among its stations.

Blocking a Bridge Port to a shared medium does not prevent that medium's stations from communicating with each other.

The access point must relay traffic among the stations, so the emulated shared medium's stations can communicate.

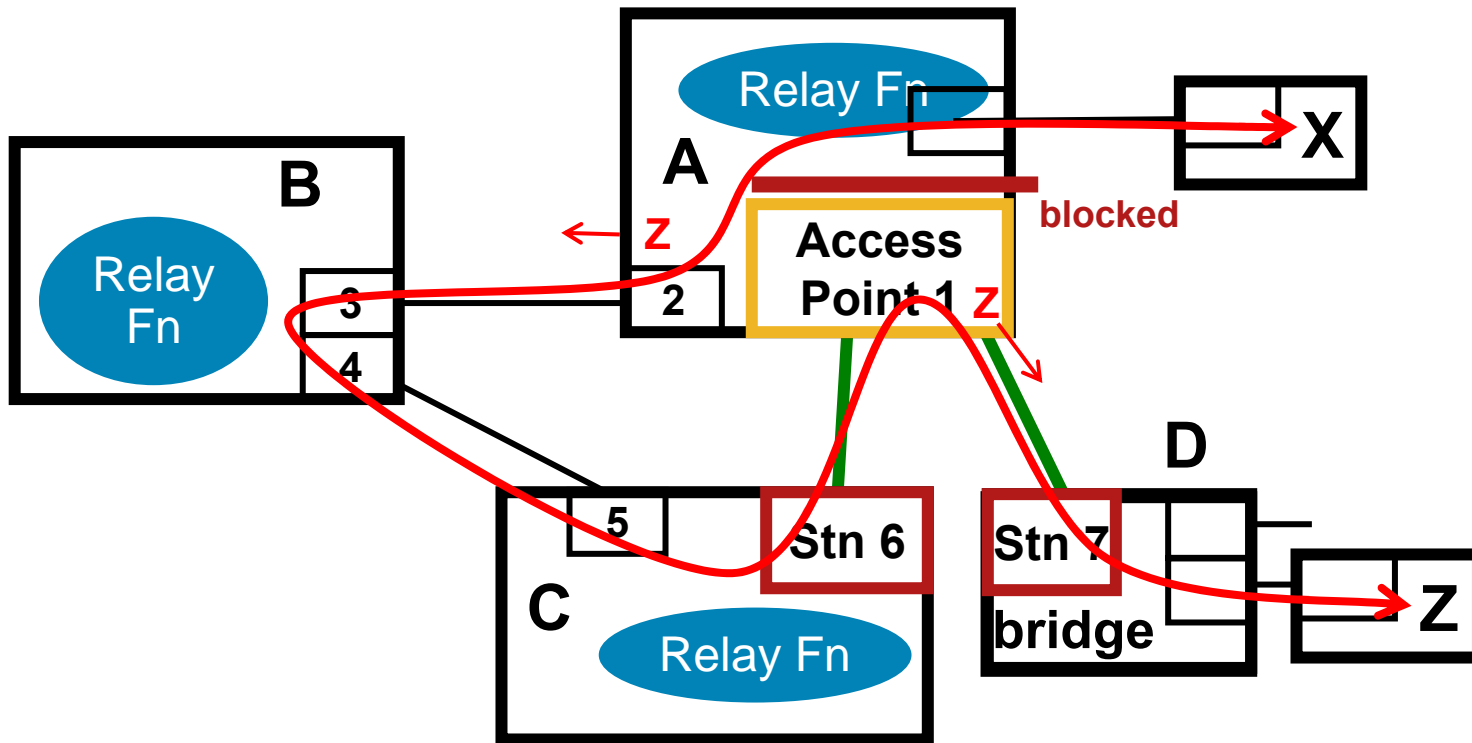
Nothing, however, can be relayed between the wired and wireless sides of the access point.

Combined Mode

Blocked Bridge Ports: Learning and Forwarding

- On the one hand, it is desirable to continue learning MAC addresses, since multiple stations can be bridges, and we want to forward frames efficiently.
- However, a bridge cannot learn MAC addresses from a blocked Bridge Port, or it will try to forward data from a wired Ethernet port to the blocked wireless port, and blackhole it, when another path through the network is available.

Combined Mode Blocked Bridge Ports



- For example, in this case, the bridge part of A learns that station Z is on port 2, while the access point part of A learns that Z is on the wireless link to D.

Combined Mode

Blocked Bridge Ports: Learning and Forwarding

- Unless it simply floods all wireless data, the blocked access point must maintain an **independent Filtering Database (FDB)**, **separate** from the main Bridge's FDB, and separate from any other blocked access point's FDB, and **perform address learning and forwarding on its own FDB** while the Bridge Port is blocked.
- If there is only one wireless medium, as for a simple 802.11b access point, this extra FDB can be a single bit added to the normal FDB.
- For access points managing multiple wireless media, the extra FDB ID information is larger, as different media may be different Bridge Ports, and hence require separate FDB IDs.

Separated Mode / Combined Mode Summary

- **Separated Mode**, in which the access point bridge's links to station bridges can be blocked separately, requires a new method to determine for which of the 2^n possible subsets of n stations a given frame is intended.

This would necessitate a change to the standards.

- **Combined Mode**, in which the access point and its associated stations act as a shared medium LAN, requires that the access point have separate filtering databases for the bridge part and the access point part.

This requires an explanation in the standards.



Other Control Protocols

Spanning Tree Protocol (STP, RSTP, MSTP)

- A given access point bridge can operate xSTP in either Combined Mode or Separated Mode; the stations do not care which mode the access point uses.
- In Separated Mode, the simulated point-to-point link to each station is a Bridge Port, and the access point bridge transmits the BPDUs separately to each station, rather than broadcasting them.
- In Combined Mode, the access point keeps a single spanning tree state for the whole emulated LAN, and must relay each BPDU received from a station to the other stations.
- The Separated/Combined Mode choice made for xSTP drives the choice for the filtering database, as well.

Multiple Registration Protocol (MRP)

- The access point bridge can operate MRP in either Combined Mode or Separated Mode; the stations do not care which mode the access point uses.
- In Separated Mode, MRP keeps a separate database in the access point bridge for each of the stations, and transmits the MRP PDUs separately to each, rather than broadcasting them.
- In Combined Mode, the access point keeps a single database for the whole emulated LAN, and must relay each MRP PDU received from a station to the other stations.
- MRP can operate in Separated Mode if xSTP is in Combined Mode, but if xSTP is in Separated Mode, so must be MRP.



Wireless links to other access point bridges

Wireless links among access point bridges

- If a number of access points are all in mutual communication with each other, they could establish an emulated LAN in the sense of “Point-to-Multipoint Bridging”,
<http://www.ieee802.org/1/files/public/docs2007/avb-nfinn-point-to-multipoint-bridging-061307.pdf>.
- However, it is difficult to guarantee the mutual reachability that the emulated LAN approach requires.
- It is much easier to treat each wireless link to another access point bridge as being equivalent to a wired point-to-point link. Then, the usual spanning tree rules can apply.

Wireless links among access point bridges

- For the sake of bandwidth efficiency, it is desirable to optimize the transmission of broadcasts, unicasts, and floods.
- Any point-to-point wireless link between access point bridges can be blocked by the spanning tree protocol.
- The problem of distributing multicasts and flooded unicasts efficiently (i.e., transmit each exactly once) is the same as the distribution problem described for Separated Mode access point bridges; the same solution should work.



There is hope

There is hope!

- As has been discussed in other documents, past and upcoming, the 802.11 four-address format for frames on the wireless medium goes a long way towards solving all of these issues, and enables wireless, wired, and mixed media bridges to work properly.

The use of this format has not, however, been standardized.

Using this format may or may not have compatibility issues with existing equipment. This is To Be Determined.

Other solutions may be found. This is To Be Determined.

There is hope!

- These problems can be solved by inventing a new device with new protocols that is similar to a bridge, but is tailored strictly for this situation and this market.
That is not what has made IEEE 802 devices successful!
- These problems can be solved by inventing new application layer protocols that locks them to end-to-end Ethernet bridged networks.
This is not what has made the Internet Protocol successful!

There is hope!

- These problems can be solved by initiating a good-faith dialog between 802.1 and 802.11 to make bridges and stations, whether wireless or wired, work together seamlessly.

They are not bridge problems.

They are not wireless problems.

They are IEEE 802 problems, and must be solved!



CISCO