



# Enhanced VRRP Support

Peter Jones  
David James  
15 May 2003

# VRRP Background

- RFC 2338 Virtual Router Redundancy Protocol provides layer 3 resiliency by allowing 2 or more routers to act as a single “virtual router”.
  - One becomes “Master”, others are “Backup” .
  - “Clients” don’t need to know that they are talking to a virtual router.
- Master uses special MAC address (VRMAC) assigned from IANA - 00-00-5E-00-01- $\{VRID\}$ 
  - $\{VRID\}$  is the VRRP Virtual Router Identifier allowing up to 255 VRRP routers on a LAN.

# VRRP Background (cont.)

- Advertisements use IP multicast packets
- The Master sends advertisements which are sent every 1 second (default) and timeout after 3 intervals.
- The Master responds to ARPs with the VRMAC in the ARP contents.
- The Traffic sent to the Virtual Router uses the VRMAC as the DA

# VRRP Background (cont.)

- Most traffic sent by the Virtual Router (i.e. all the data traffic) uses it's real MAC as the SA.
- VRMAC address is used as the source MAC address in VRRP messages sent by the Master to enable bridges to learn the VR location.

# VRRP and Draft 2.2

- VRRP can function correctly with draft 2.2 of 802.17, but not efficiently
- All transmissions to the VRMAC will be flooded as the VRMAC is not in the topology image, so these frames will not get spatial reuse.
- A client wishing to act as a VRRP router will have to use “Bridge Mode” to receive transmissions to the VRMAC.
- This client will need to filter out other traffic such as that being sent between bridges on the ring.

# VRRP and Draft 2.2 (cont.)

- D2.2 does not discuss how to implement VRRP over RPR and it's not obvious at first reading.
- It can be made to work, but it's not efficient
- Are we happy to have all transmission to the VRRP Master be flooded?
- Does this fit RPR's position in the market

# Possible Options

- We could
  - Do nothing
  - Add some text into the draft to explain how VRRP and RPR work together
  - Make changes to the current behavior to support VRRP and spatial reuse.
- The rest of this presentation discusses a proposal to add support for VRRP and Spatial reuse to 802.17
- This proposal has been discussed in the BAH and PAH adhocs, although this presentation is not a BAH or PAH agreed position.

# Enhanced VRRP support

- Add the concept of “secondary addresses” to RPR topology and client interface data path
  - Registered by the client when it becomes a VRRP master
  - Advertised through topology TLV message
  - Becomes part of the topology image as extra columns in the table
  - VRMAC DA address is processed in the ringlet selection block in frame transmission to:
    - Determine the associated station
    - Encapsulate in the extended frame format to transmit the frame to the station’s real MAC address
  - Check the extended DA address on the copy to client receive processing on the destination station

# Enhanced VRRP support (cont.)

- Impact of this proposal
  - New TLV attribute advertised and processed in topology
  - Secondary entries processed in ringlet selection (associative lookup??)
  - Extended frame format used without flooding
  - Validation check in copy to client when stripping extended frame header
  - NOT Affected
    - Fairness
    - Protection
    - Transit path

# Enhanced VRRP support (cont.)

- Limit total number of active secondary addresses on the ring to control maximum size of table in ringlet selection
  - 16 or 32 sounds like a good number.
- Need to support at least two secondary addresses per station.
  - Allows for load sharing between two real routers supporting two virtual routers
  - Could limit to maximum 2 per station if this minimizes the impact to the topology protocol and data structures
- Extra fields in the MIB to report the active secondary entries in the topology table
- Need to have defined the TLV update/change processing properly

# Conclusions

- At a minimum, we need to add text to the draft to explain how VRRP and RPR will interact.
- I recommend we make changes to D2.2 as specified in comment numbers 65, 368, 468, 644 and outlined in Peter\_Jones\_vrrp\_proposal\_01.pdf to permit spatial reuse for frames sent to VRRP.