Spatially aware sublayer support of VRRP

Marc Holness, Nortel Networks

IEEE 802.17 WG — 802.17b SG

San Antonio, Texas

November, 2004

Agenda

- Objectives
- Problem overview
- Solution overview
- Spatially aware sublayer operations

Terminology and terms

- Directed transmissions – Refers to a RPR source station transmitting to a designated (unicast) destination address on the ring
- Undirected transmission – Refers to a RPR source station flooding a frame over the ring
- Remote address – A MAC address that is not local to the ring (i.e., not a local RPR station MAC address)
- SAS – Spatially aware sublayer
- VRRP – Virtual router redundancy protocol (as specified in RFC 3768)
- VMAC – Virtual MAC (or VRRP group)

Objectives

- Outline how the spatially aware sublayer (SAS) can support VRRP applications

Problem overview

Typical WAN

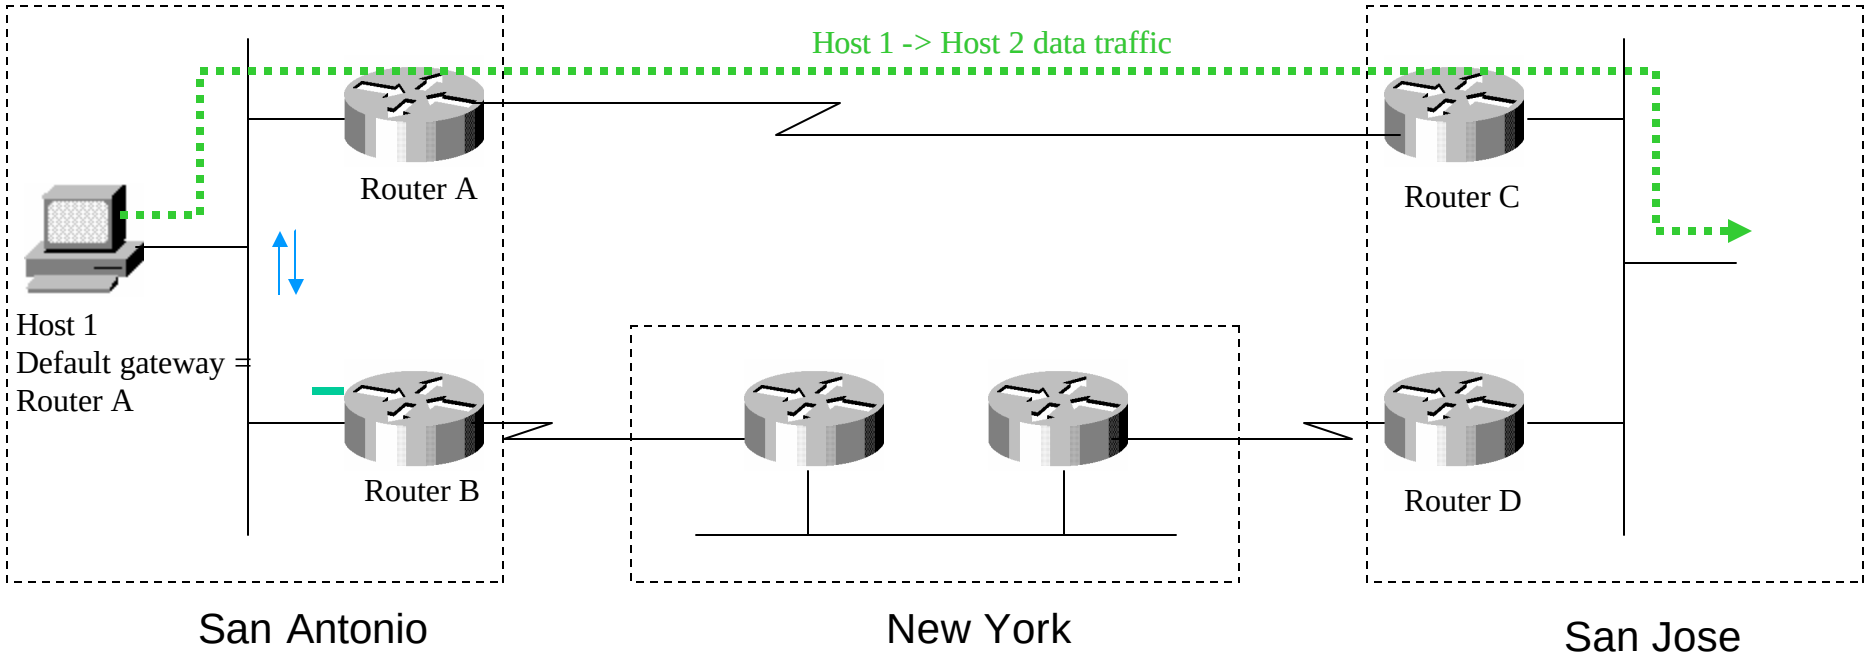

Host 1 uses default gateway Router A to get to Host 2

Problem overview

Typical WAN

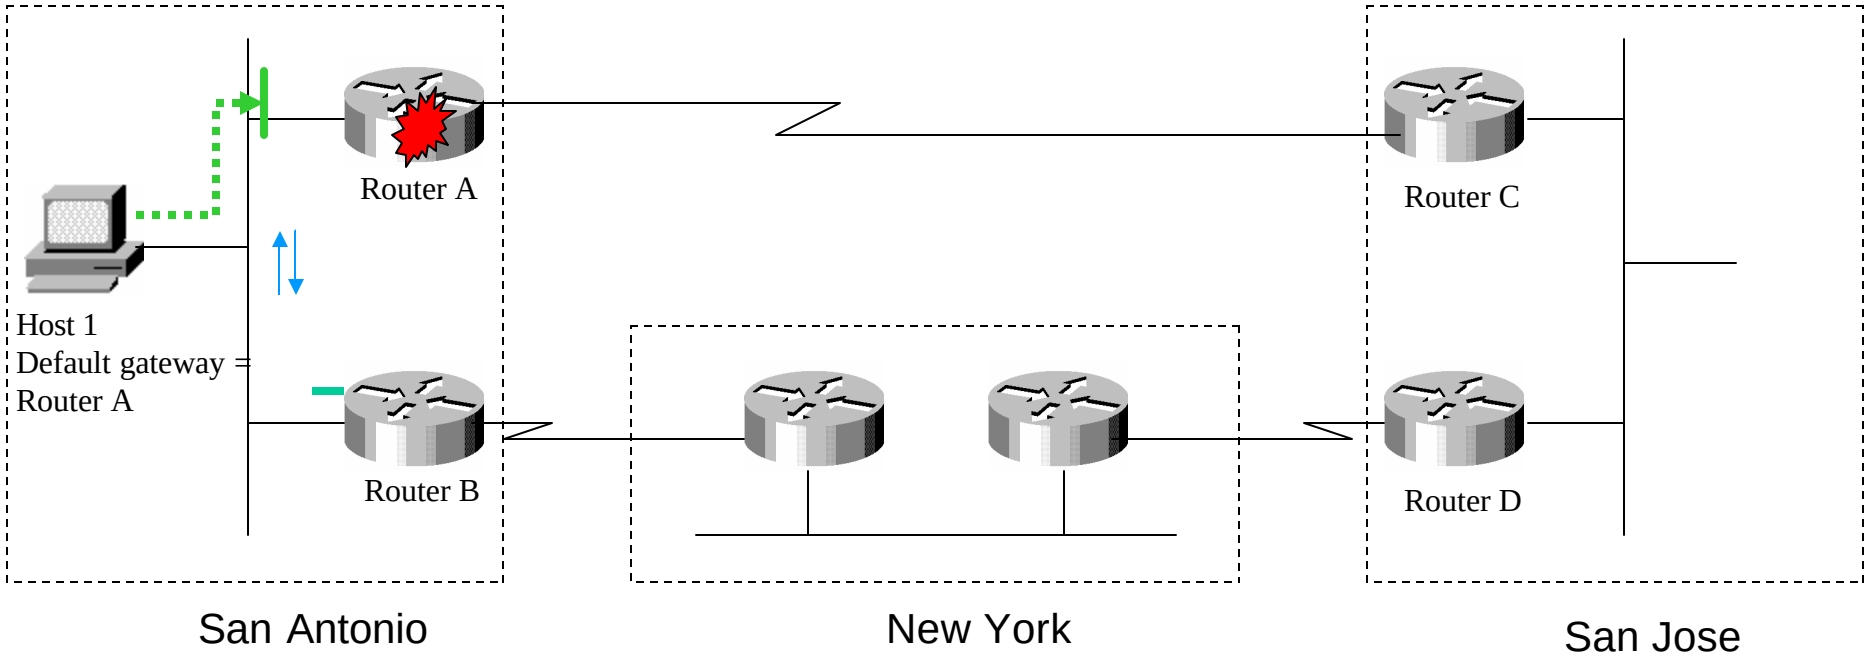

If fault occurs at Router A (or link to Router A), communication from Host 1 to Host 2 is interrupted

Solution overview

- A solution that allows for fast convergence and does not affect current host implementations is required

VRRP is born!

VRRP basics

Physical view

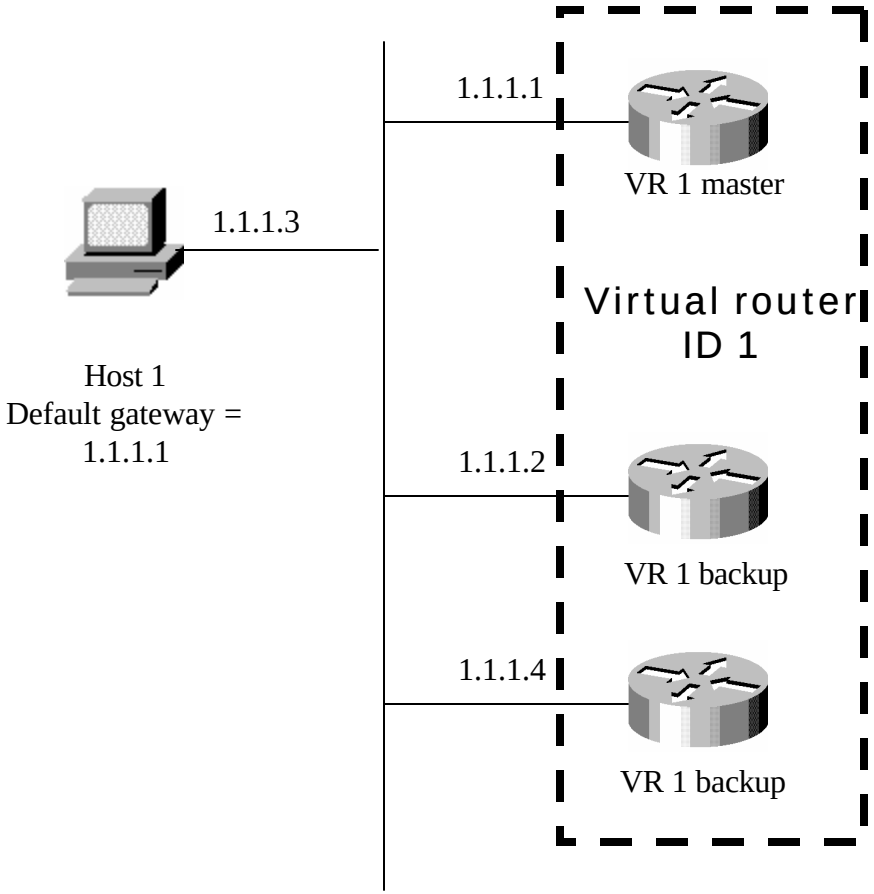

Logical view

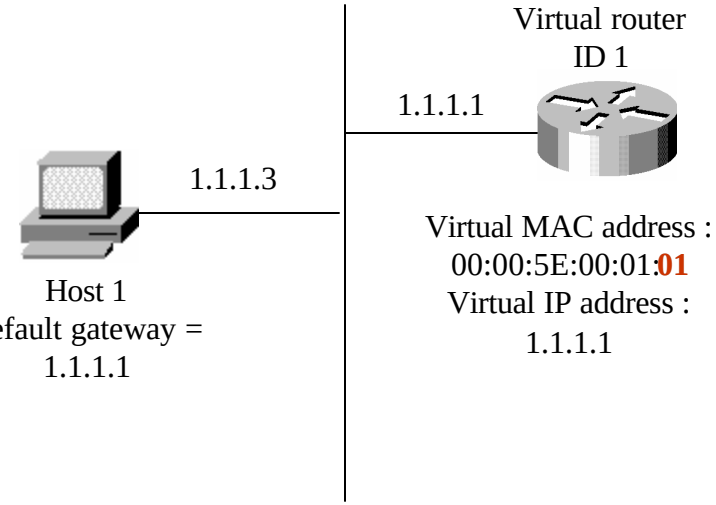

Virtual Router ID 1
Virtual MAC address - 00:00:5E:00:01:**01**
Virtual IP address - 1.1.1.1

VRRP most common implementation

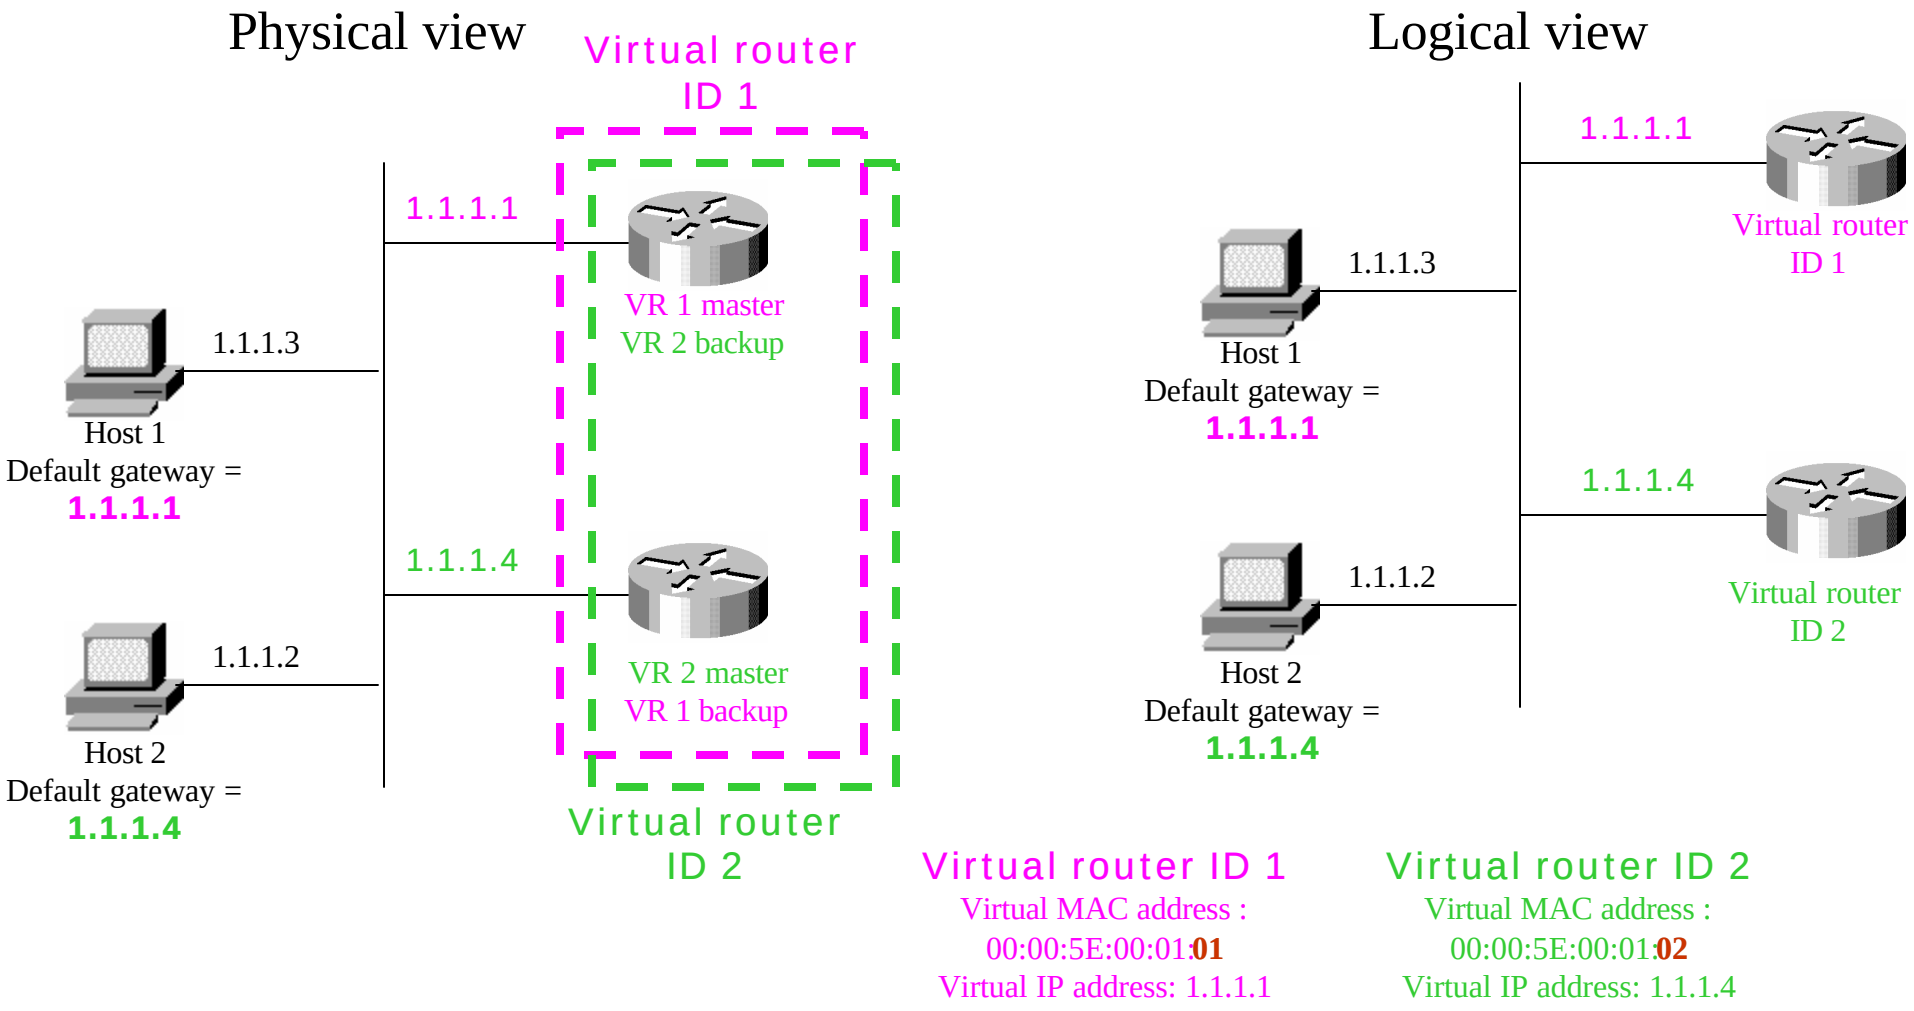

VRRP operation

Normal operation

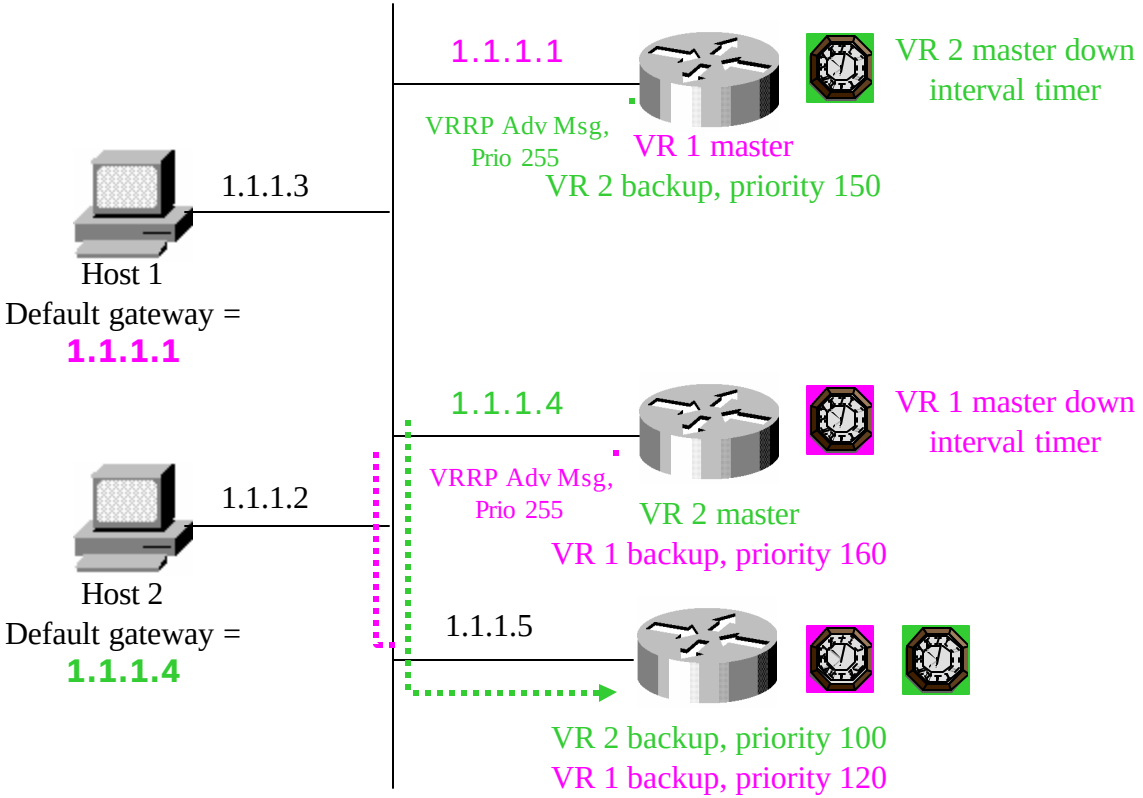

VRRP operation

Failure operation (master down interval time out)

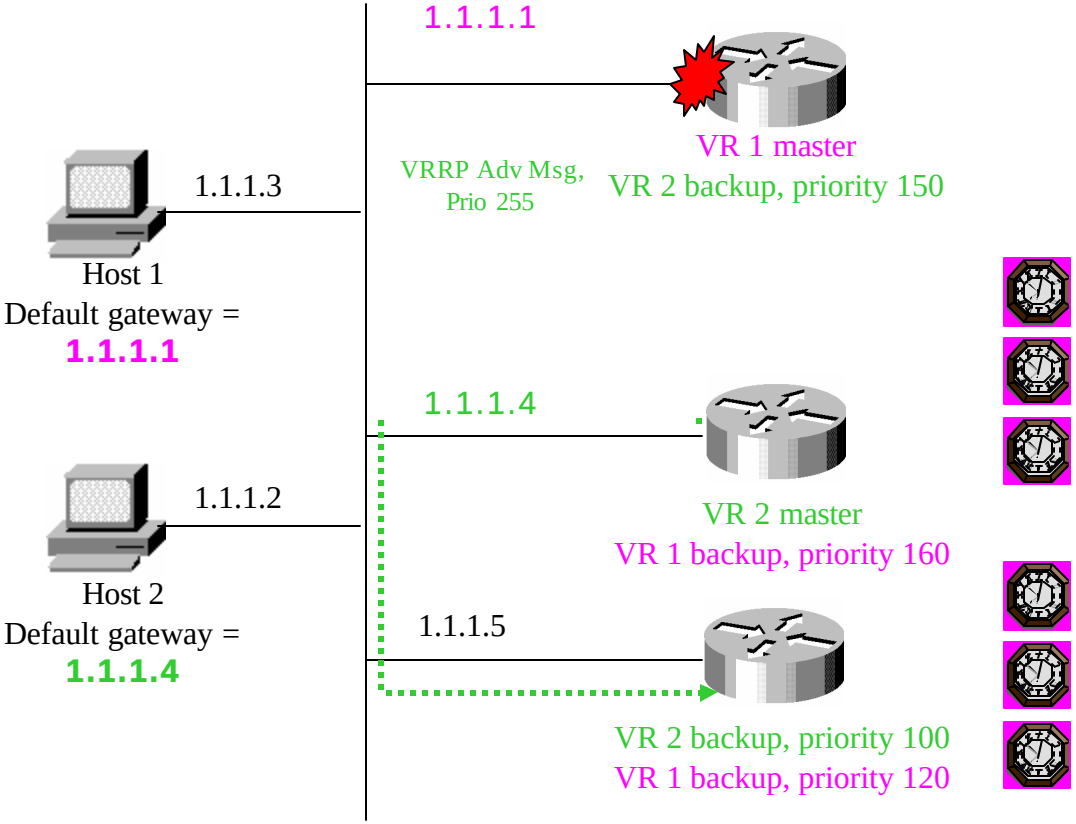

VRRP operation

Failure operation (election process)

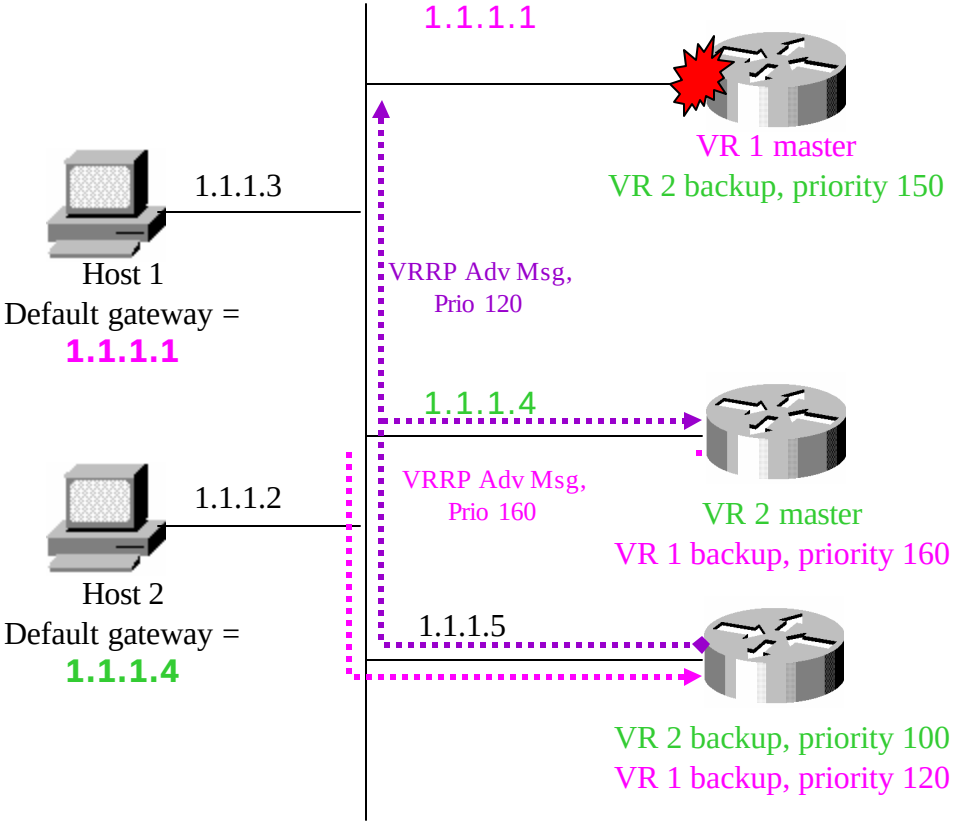

VRRP operation

Failure operation - and the winner is ...

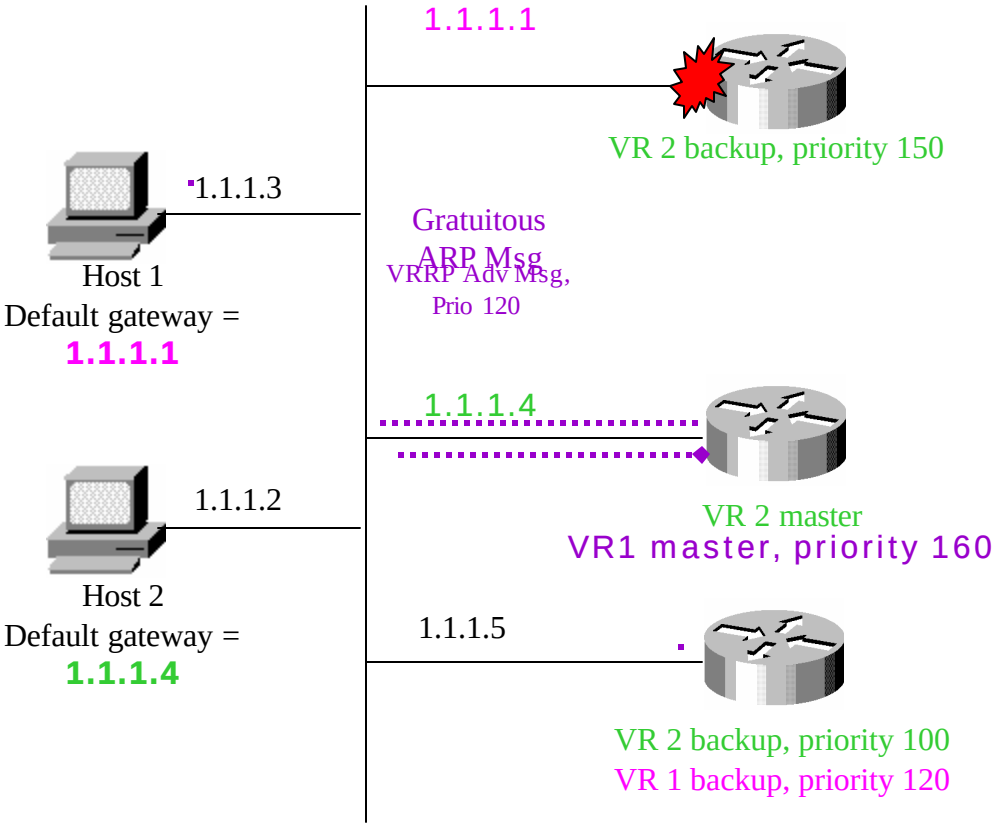

VRRP advertisement message

- Sent periodically by master router in virtual router group (i.e., VRID)
  - Received and processed by backup routers in virtual router group

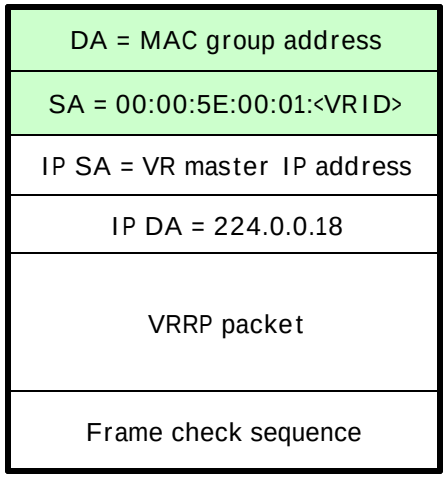

Gratuitous ARP message

- Sent by master router of virtual router group (i.e., VRID) at startup and when master router fail over occurs
  - Host devices receiving gratuitous ARP update ARP cache tables
  - Intermediate bridging devices, if present, will update FDB entries associated with virtual router master interface MAC address and received port

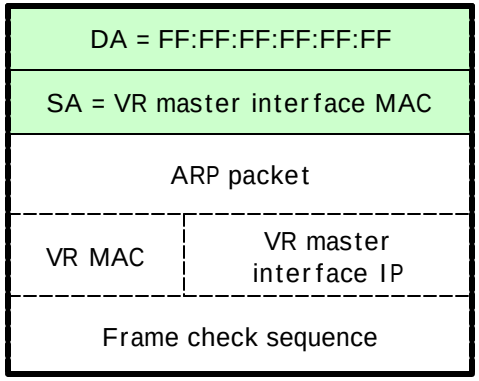

Host ARP request message

- Sent by host at startup time to discover MAC address of virtual router master
  - Intermediate bridging devices, if present, will update FDB entries associated with host MAC address and received port

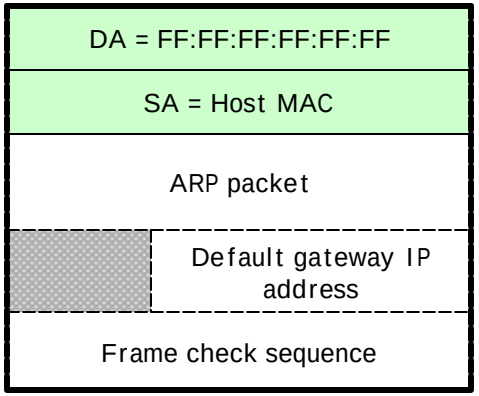

Host ARP reply message

- Sent by virtual router master in response to host ARP request
  - Intermediate bridging devices, if present, will update FDB entries associated with virtual router interface MAC address and received port

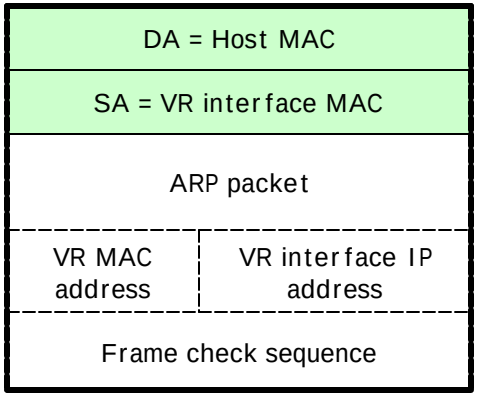

RPR and VRRP network references (1)

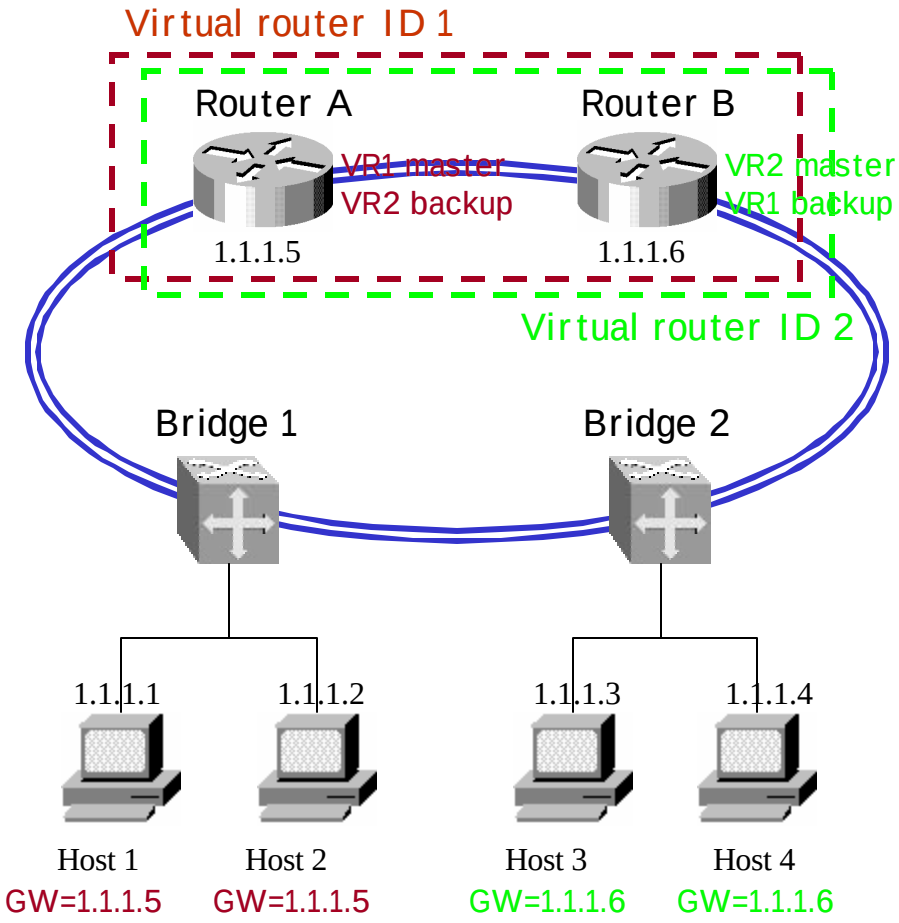

RPR and VRRP network references (2)

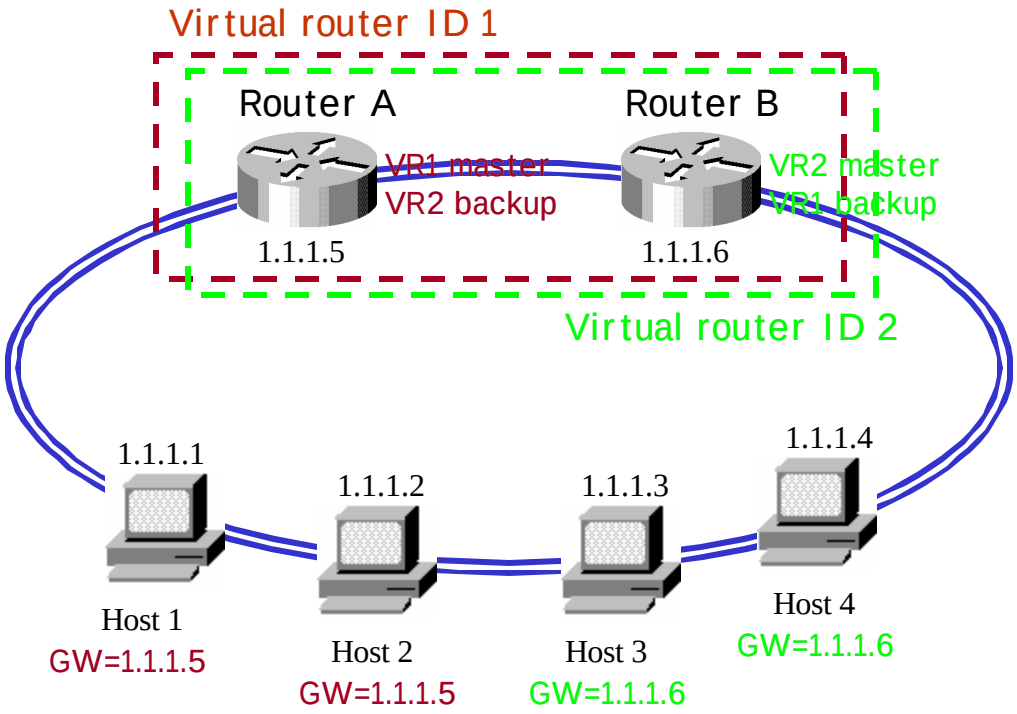

RPR and VRRP network references (3)

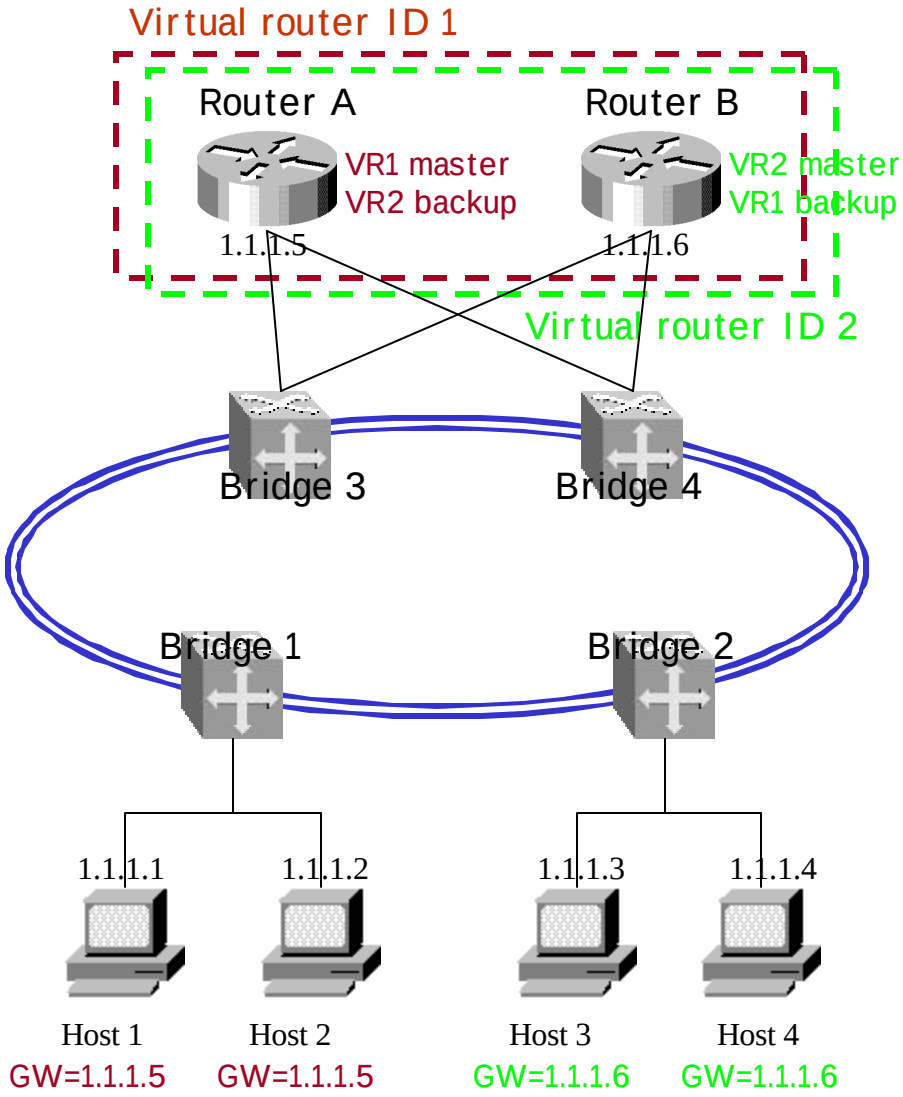

802.17 support of VRRP

- Virtual MAC (i.e., 00:00:5E:00:01:<VRID>) used by VRRP does not represent an address associated with a RPR station
  - Consequently 802.17 MAC will, by default, flood frames with this MAC address
  - Results in lack of spatial reuse and poor bandwidth utilization over the ring

802.17b SAS interactions with VRRP

SAS support of VRRP

- Advertisement messages over 802.17b

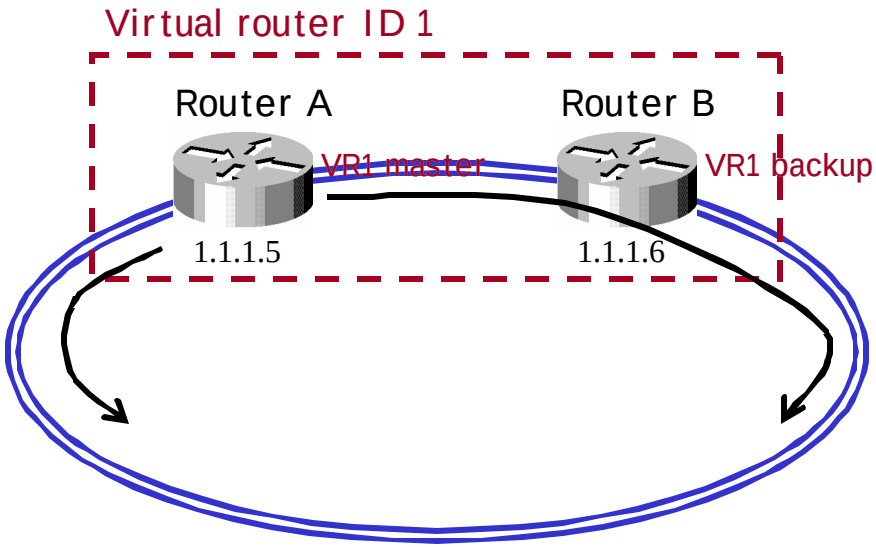

SAS

- Associate VMAC with RPR station MAC attachment
- Frames transmitted to VMAC (which is considered a remote address) can now be associated with an RPR station MAC, resulting in directed transmissions and achieving spatial reuse over RPR

SAS support of VRRP

- Gratuitous ARP messages over 802.17b

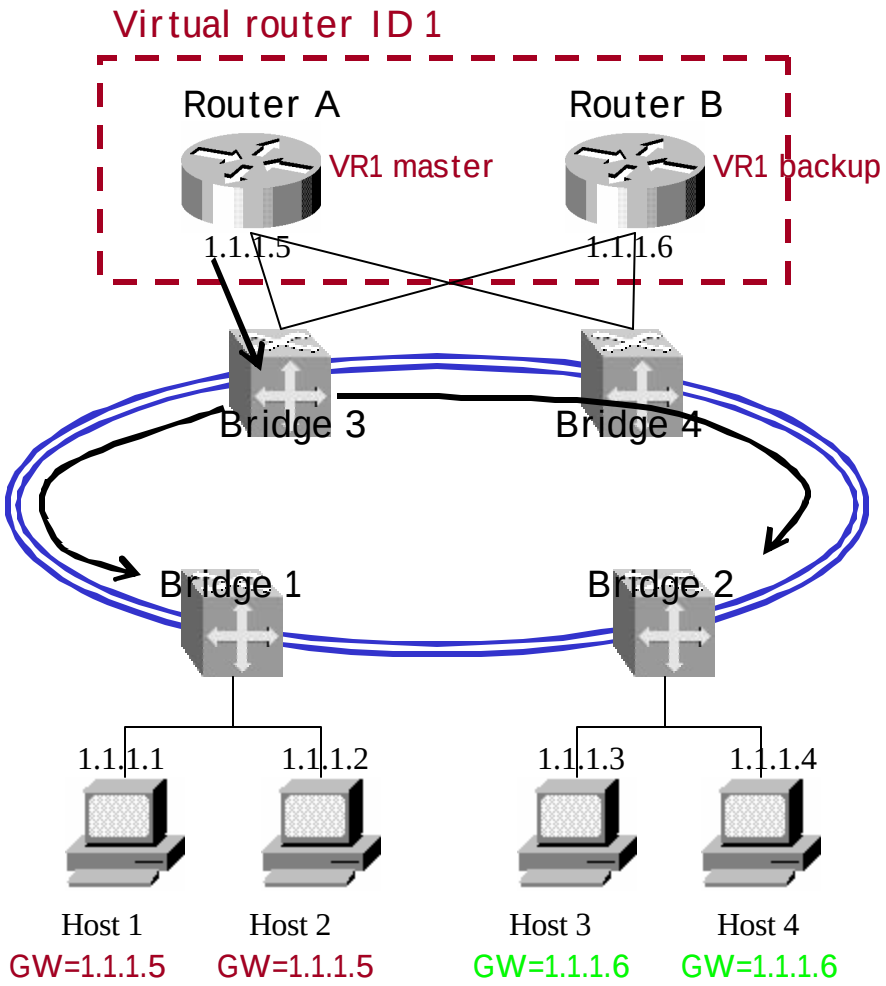

SAS

- Associate VR master interface MAC (remote address) with RPR station MAC attachment
- Frames transmitted to the VR master interface MAC (remote address) can now be associated with an RPR station MAC, resulting in directed transmissions and achieving spatial reuse over RPR

SAS support of VRRP

- Host ARP request messages over 802.17b

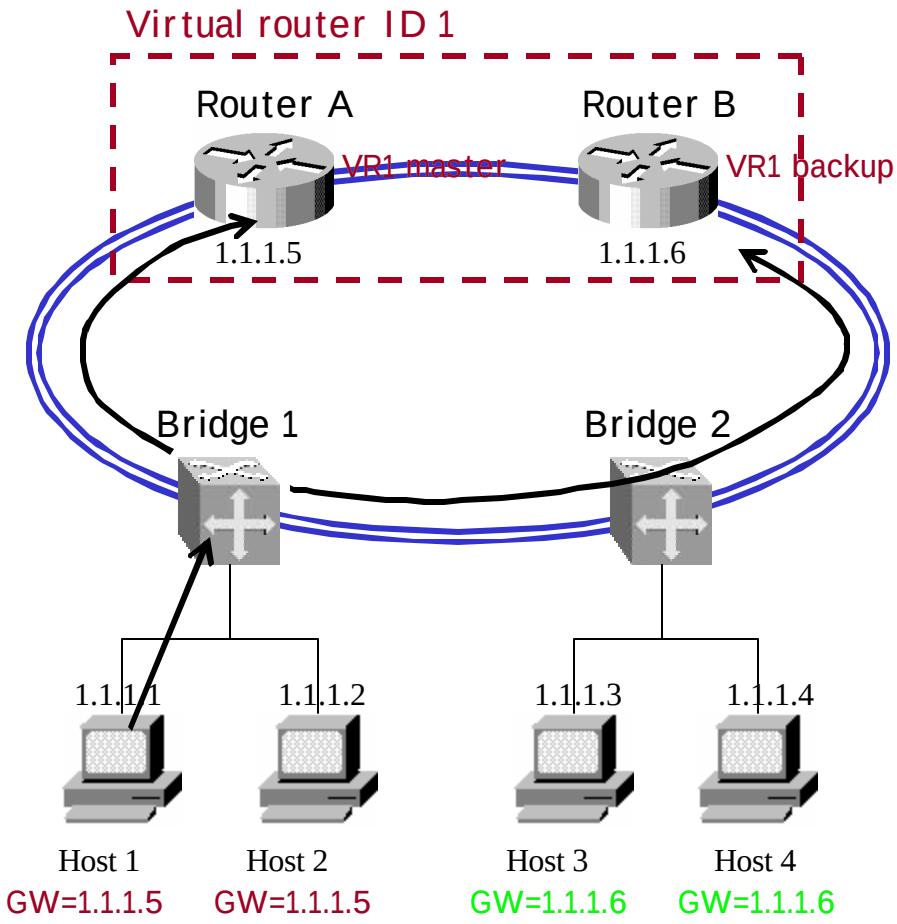

SAS

- Associate host MAC address (remote address) with RPR station MAC attachment
- Frames transmitted to the host MAC address (remote address) can now be associated with an RPR station MAC, resulting in directed transmissions and achieving spatial reuse over RPR

SAS support of VRRP

- Host ARP reply messages over 802.17b

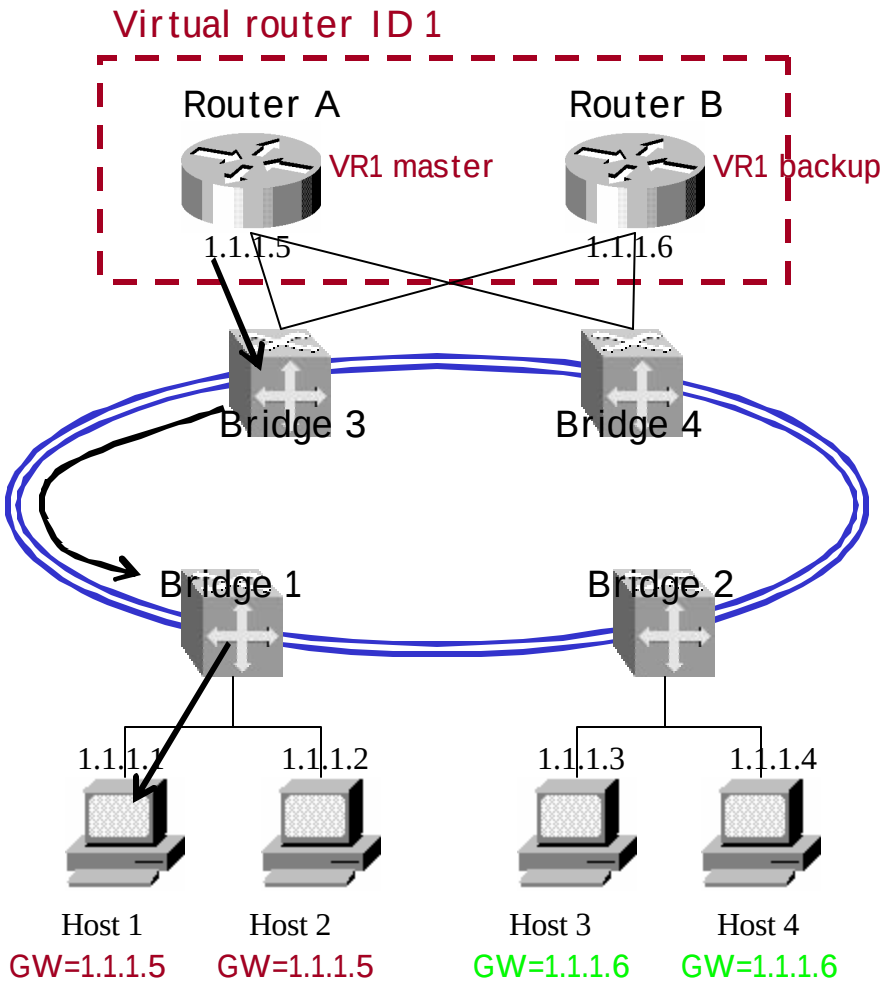

SAS

- Associate VR master interface MAC (remote address) with RPR station MAC attachment
- Frames transmitted to the VR master interface MAC (remote address) can now be associated with an RPR station MAC, resulting in directed transmissions and achieving spatial reuse over RPR

SAS support of VRRP

- VRRP normal operation

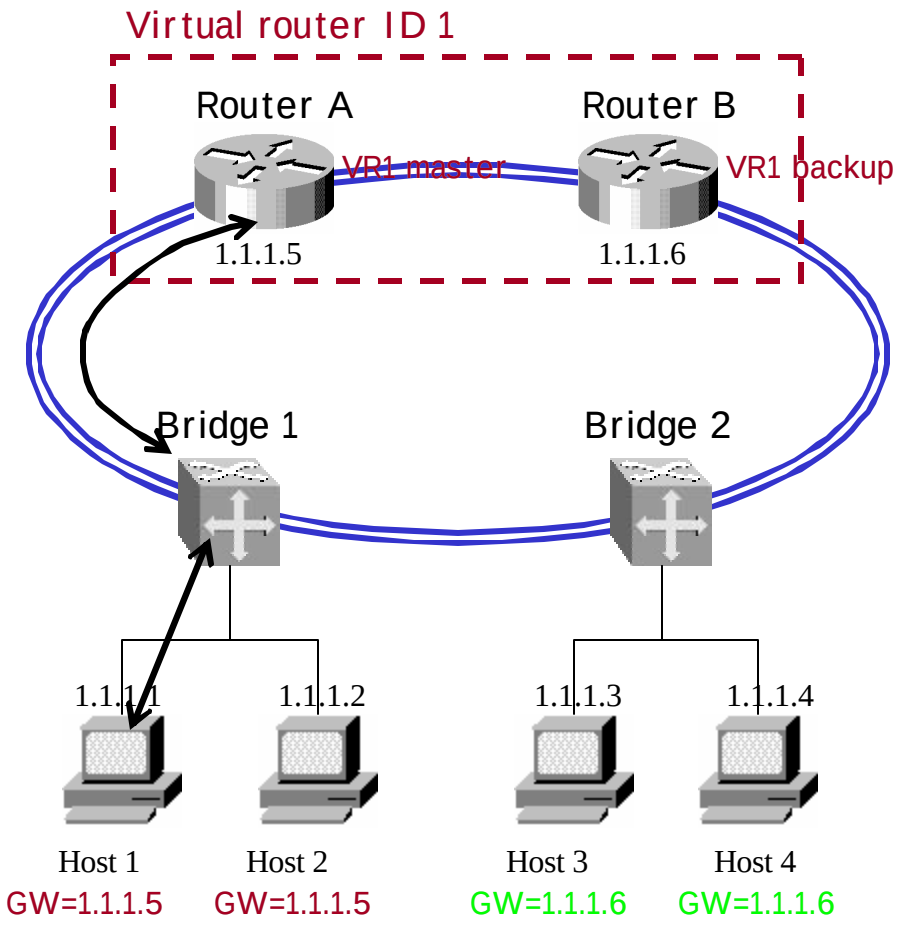

SAS

- Data transmissions from Host 1 to off-net host via VR master router will result in SAS DB lookup of VMAC
  - RPR station MAC entry found in SAS DB will result in directed transmissions and thus achieve spatial reuse over RPR
- Data transmissions from VR master router to Host 1 will result in SAS DB lookup of Host 1 MAC address (remote)
  - RPR station MAC entry found in SAS DB will result in directed transmissions and thus achieve spatial reuse over RPR

Summary

- SAS supports VRRP applications over RPR
  - Provides spatial reuse improvements
  - No additional RPR configuration required
  - VMAC to RPR station MAC associations are auto-learned
  - Supports various VRRP network configurations involving RPR, routers, bridges, and hosts

Backup

VRRP/RPR packet walk-thru

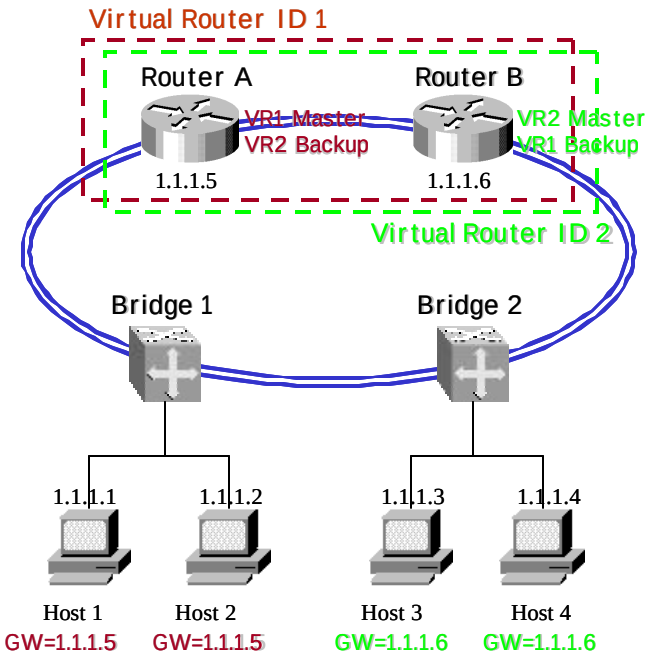

VRRP startup

1. VR1 master (RouterA) sends advertisement message
   - VR1 backup (RouterB) receives advertisement message and resets Master_Down_Timer
   - SAS in RPR RouterB, RPR Bridge1 & RPR Bridge2 associates VMAC1 with RouterA MAC (MAC_{RA})
   - Bridge1 forwards VRRP advertisement message to Host1 and Host2
   - VR1 master (RouterA) sends advertisement message every Advertisement_Interval (e.g., 1 sec)
2. VR1 master (RouterA) broadcast gratuitous ARP message
   - SAS in RPR RouterB, RPR Bridge1 & RPR Bridge2 associates VR master interface MAC with RouterA MAC (MAC_{RA})
   - Bridge1 forwards gratuitous ARP to Host1 and Host2
   - Host1 and Host2 updates ARP cache table

VRRP/RPR packet walk-thru

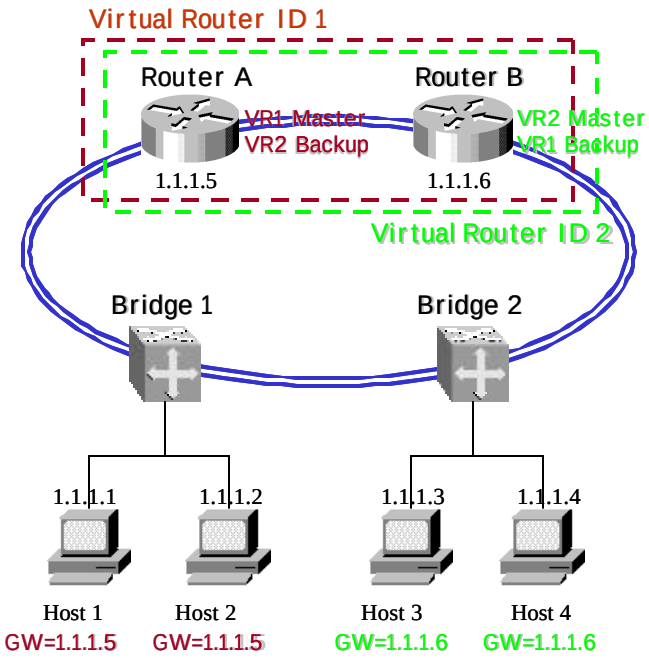

Host startup

3. Host1 broadcast ARP request message
   - RPR Bridge1 floods frame over RPR
   - SAS in RPR RouterA & RouterB, RPR Bridge1 & RPR Bridge2 associates Host1 MAC (MAC_{H1}) with RPR Bridge1 MAC (MAC_{B1})
   - VR1 master (RouterA) receives ARP request and processes message
4. VR1 master (RouterA) responds to ARP request with ARP reply message
   - Host1 updates ARP cache table

VRRP/RPR packet walk-thru

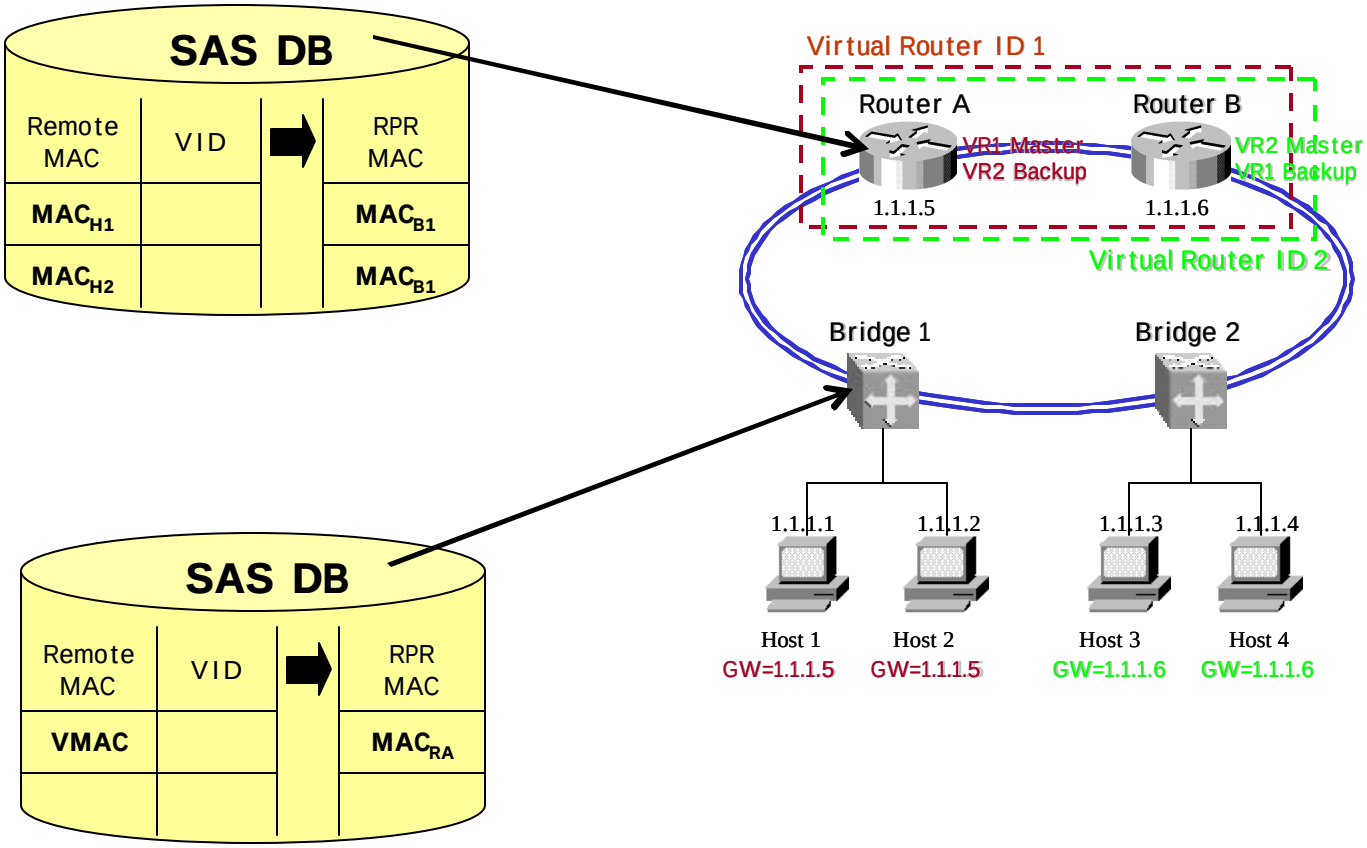

After VRRP initialization signaling and Host ARP process, the SAS data bases are updated.

VRRP/RPR packet walk-thru

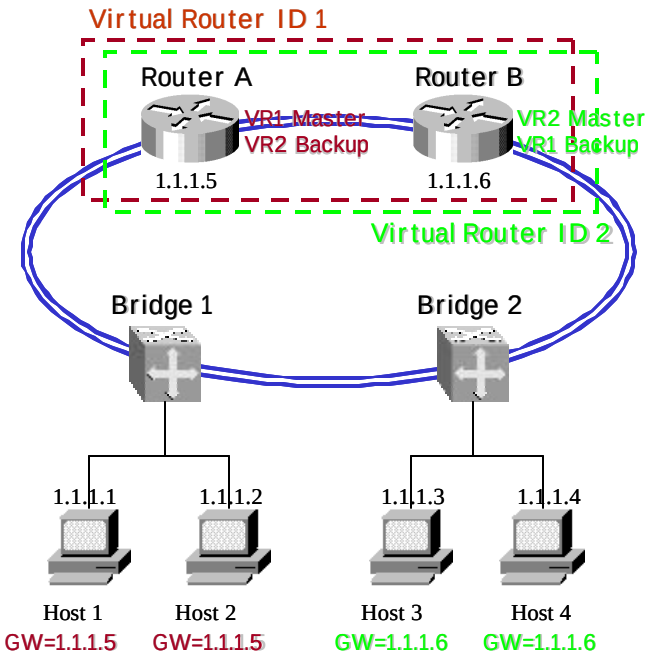

Normal operation

5. Host1 sends message to off-net Host via VR1 (using VMAC in destination address)
   - RPR Bridge1 MAC indexes SAS DB using VMAC and discovers RPR MAC address MAC_{RA} is RPR destination target address to perform directed transmission to RPR station RouterA
6. Off-net Host messages Host1 via VR1 master (RouterA)
   - RPR RouterA MAC indexes SAS DB using MAC_{H1} and discovers RPR MAC address MAC_{B1} is RPR destination target address to perform directed transmission to RPR station Bridge1
   - Bridge1 dispatches packet to Host1

SAS provides spatial reuse over RPR involving VRRP interactions.

VRRP/RPR packet walk-thru

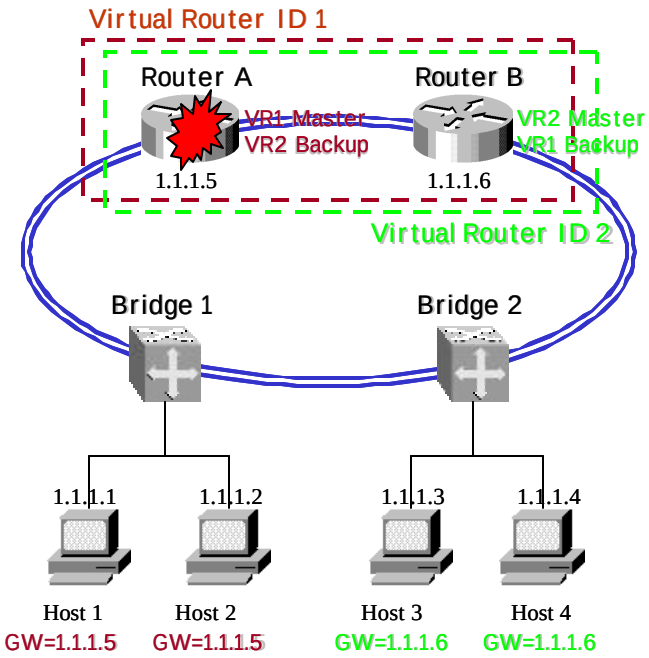

VR1 master (RouterA) failure

7. VR1 backup routers (RouterB) master down interval timeout occurs; VR1 master router election process occurs, resulting in RouterB becoming the new VR1 master
8. VR1 master (RouterB) sends advertisement message
   - SAS in RPR Bridge1 & RPR Bridge2 associates VMAC1 with RouterB MAC (MAC_{RB})
   - Bridge1 forwards VRRP advertisement message to Host1 and Host2
   - VR1 master (RouterB) sends advertisement message every Advertisement_Interval (e.g., 1 sec)
9. VR1 master (RouterB) broadcast gratuitous ARP message
   - SAS in RPR Bridge1 & RPR Bridge2 associates VR master interface MAC with RouterB MAC (MAC_{RB})
   - Bridge1 forwards gratuitous ARP to Host1 and Host2
   - Host1 and Host2 updates ARP cache table

VRRP/RPR packet walk-thru

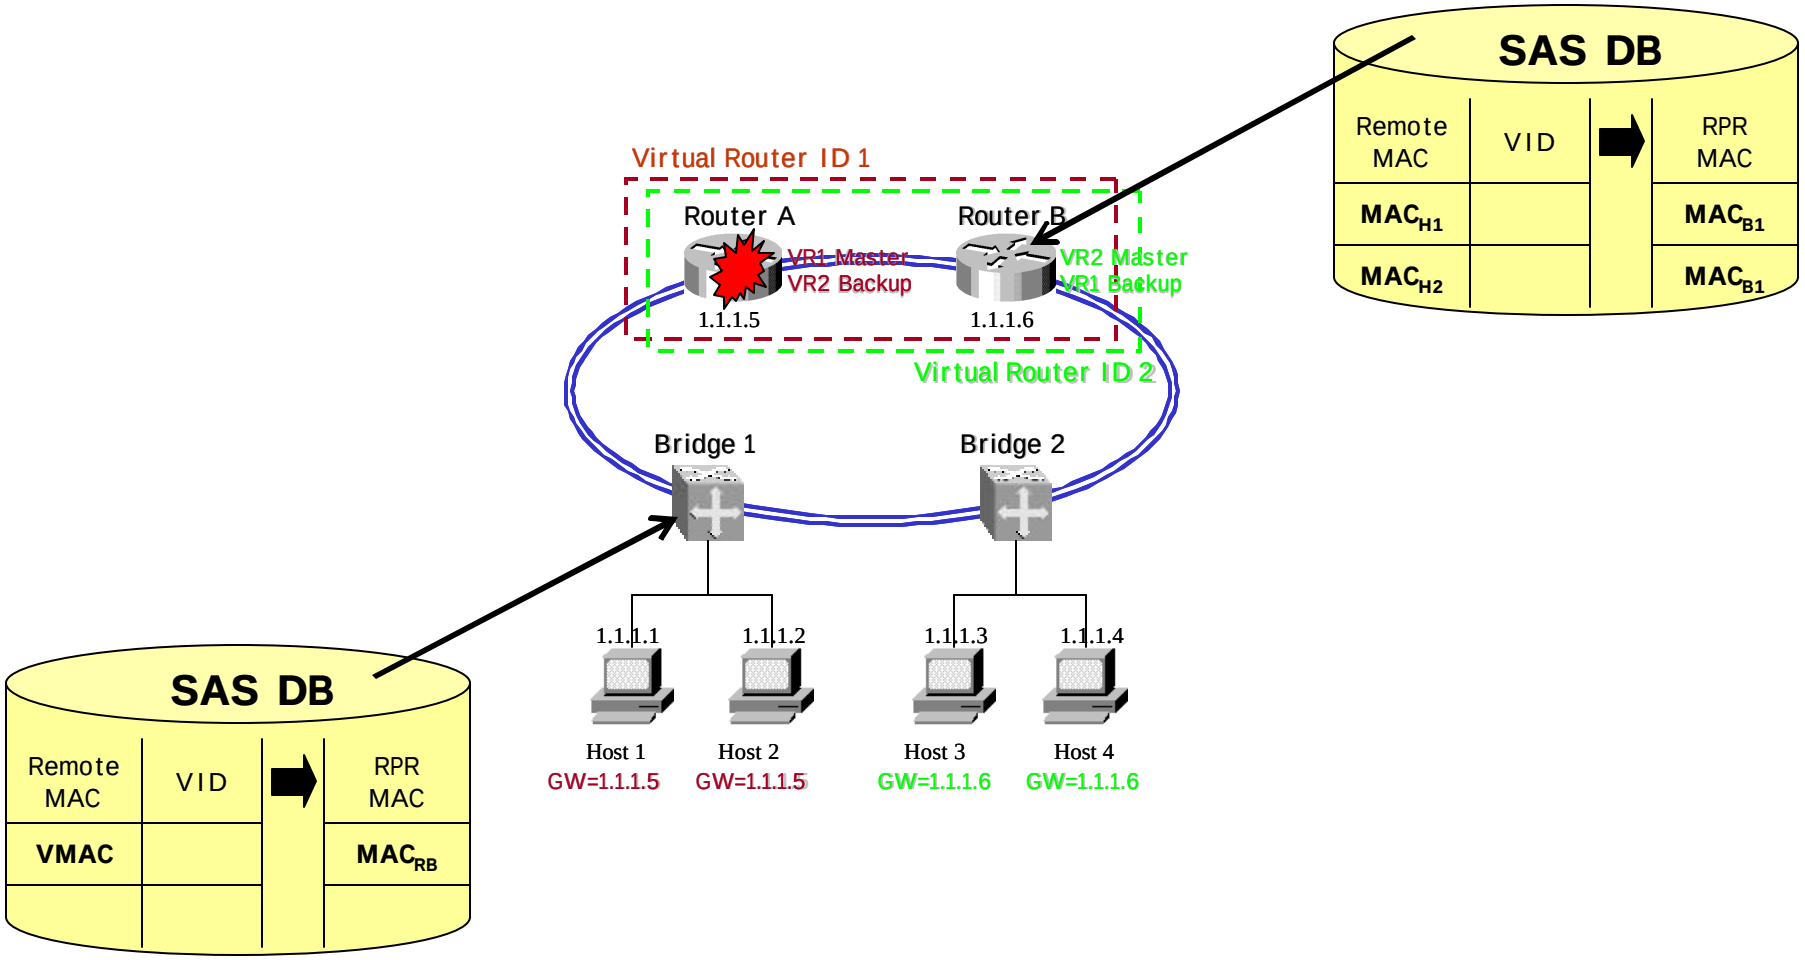

VRRP process to facilitate master router switchover results in the SAS data bases to be updated.

VRRP/RPR packet walk-thru

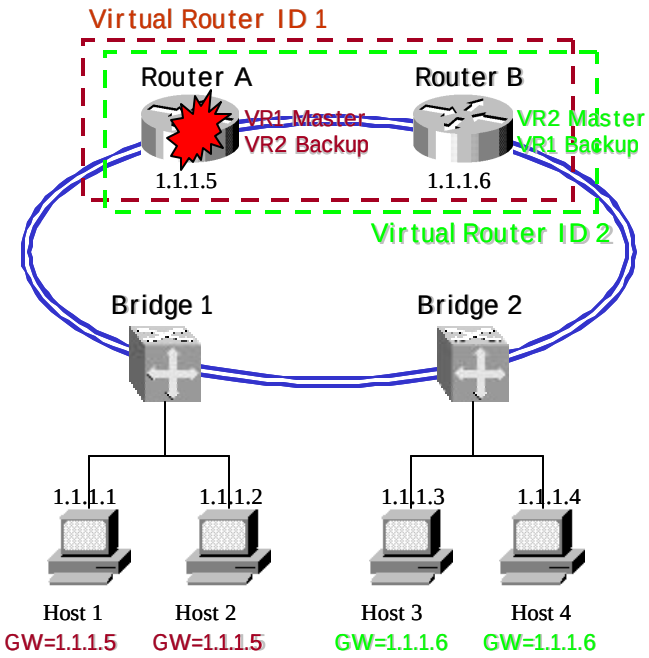

Normal operation

10. Host1 sends message to off-net Host via VR1 (using VMAC in destination address)
    - RPR Bridge1 MAC indexes SAS DB using VMAC and discovers RPR MAC address MAC_{RB} is RPR destination target address to perform directed transmission to RPR station RouterB
11. Off-net Host messages Host1 via VR1 master (RouterB)
    - RPR RouterB MAC indexes SAS DB using MAC_{H1} and discovers RPR MAC address MAC_{B1} is RPR destination target address to perform directed transmission to RPR station Bridge1
    - Bridge1 dispatches packet to Host1

SAS continues to provide spatial reuse over RPR involving VRRP interactions.