

Router Virtualization Protocols

Raj Jain

Washington University in Saint Louis

Saint Louis, MO 63130

Jain@cse.wustl.edu

These slides and audio/video recordings of this class lecture are at:

<http://www.cse.wustl.edu/~jain/cse570-23/>

Student Questions



1. Hot Standby Router Protocol (HSRP)
2. Virtual Router Redundancy Protocol (VRRP)
3. Virtual Routing and Forwarding (VRF)
4. Virtual Device Context (VDC)

Student Questions


Router Virtualization

- ❑ Multiple routers act as one router
 - Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP)
 - Virtual Routing and Forwarding (VRF)
- ❑ One router acts as multiple routers
 - Virtual Device Context (VDC)

Student Questions

Hot Standby Router Protocol (HSRP)

- ❑ Problem: What to do if the default gateway fails?
- ❑ Solution: Multiple default gateways act as one virtual default gateway.
- ❑ All routers have a real IP (RIP) and a virtual IP (VIP).
Clients send to VIP
- ❑ Only the active/master router forwards packets.
- ❑ Standby routers keep track of the active/master router.
Elect a new active router if the current active router fails.
- ❑ **Hot Standby Router Protocol (HSRP)** is described in RFC 2281.



Student Questions

Virtual Router Redundancy Protocol (VRRP)

- ❑ HSRP is Cisco proprietary.
- ❑ **Virtual Router Redundancy Protocol (VRRP)** described in RFC 5798 is the IETF standard version and is similar.
- ❑ Each VRRP is limited to a single subnet.
Does not advertise routes beyond that subnet.
Does not affect routing tables.

Student Questions

- ❑ Despite the little difference, VRRP is safer than HSRP, right?


VRRP is standard and newer.

- ❑ Is HSRP used these days? And what about VRRP?

Both are used.

Routing Information Base (RIB)

- ❑ RIB: Routing table constructed by various routing protocols, such as Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Routing Information Protocol (RIP)
- ❑ For each destination prefix, RIB entries point to which router the packet should be sent to. The router may or may not be adjacent (particularly for BGP).
- ❑ Example: R_1 's RIB will show R_4 as the path to $128.28.23.0/8$
- ❑ RIB preparation is a control plane activity.
- ❑ The control plane is implemented using general-purpose processors.




Student Questions

- ❑ Is the router pointed by RIB the virtual router or one of the physical routers inside a virtual router?

It can be any router. Virtual or physical.

Forwarding Information Base (FIB)

- ❑ FIB: The table used to forward packets. It lists the *next hop* for each destination prefix.
- ❑ FIB is used in the data plane, so it must be fast.
- ❑ The data plane is implemented using special-purpose network processors.
- ❑ FIB is constructed from RIB. When the entries are copied, the next hop is resolved, outgoing interfaces are computed, and multiple entries are created if there are multiple paths to the destination prefix.
- ❑ Example: R₁'s FIB will have 2 entries for 128.28.23.0/8 via R₂ and R₃.



Student Questions

- ❑ Why do we need both RIB and FIB? Would FIB not be sufficient?

FIB resides in high-speed memory.


- ❑ Is FIB a virtual implementation of a regular routing table?

It is the routing table.

- ❑ Will FIB find the shortest path or just constructed from all paths? For example, for R₃'s FIB for 128.28.23.0/8, will the entries be only (R₄) or (R₄ and R₁)?

Usually, all routers will have multiple entries.

Virtual Routing and Forwarding (VRF)



- ❑ Allows multiple virtual routers within a single physical router.
 - Each virtual router has its own routing table, network interfaces, and routing protocols
 - Each virtual router has its forwarding information base (FIB) and routing information base (RIB)
 - VRFs exchange routing (e.g., OSPF) information with other virtual and physical routers.
 - Users of one VRF cannot communicate with users of another VRF unless explicitly configured.

Ref: http://en.wikipedia.org/wiki/Virtual_Routing_and_Forwarding

Washington University in St. Louis


<http://www.cse.wustl.edu/~jain/cse570-23/>

©2023 Raj Jain

Student Questions

VRF (Cont)


- ❑ Originally designed for carrier networks to allow enterprises to have a virtual WAN using MPLS.
- ❑ Customer edge (CE) routers forward to provider edge (PE) routers. Entry PE routers encapsulate packets and forward them to exit PE. Exit PE decapsulates and gives it to CE.
- ❑ PE routers use a virtual routing table containing only routes to the customer's virtual PEs.
Core routers in the carrier's network are unchanged.



Student Questions

VRF Tunneling

- ❑ VRFs connected over multiple hops require tunneling. GRE and MPLS are two tunneling techniques used in WANs.
- ❑ VRF Lite = w/o MPLS = hop-by-hop
- ❑ VRF is now used in data centers. Multiple VRFs can share an interface using VLANs.




Student Questions

- ❑ How does VRF lite use hop-by-hop?
Each hop has a TCP/UDP end-point.
- ❑ Could you explain this picture in detail, please? *Sure.*

VRF Example

□ Data Center Segmentation:

Physical
Topology




Source: Santana 2014

Student Questions

VRF Example (Cont)

Logical
Topology

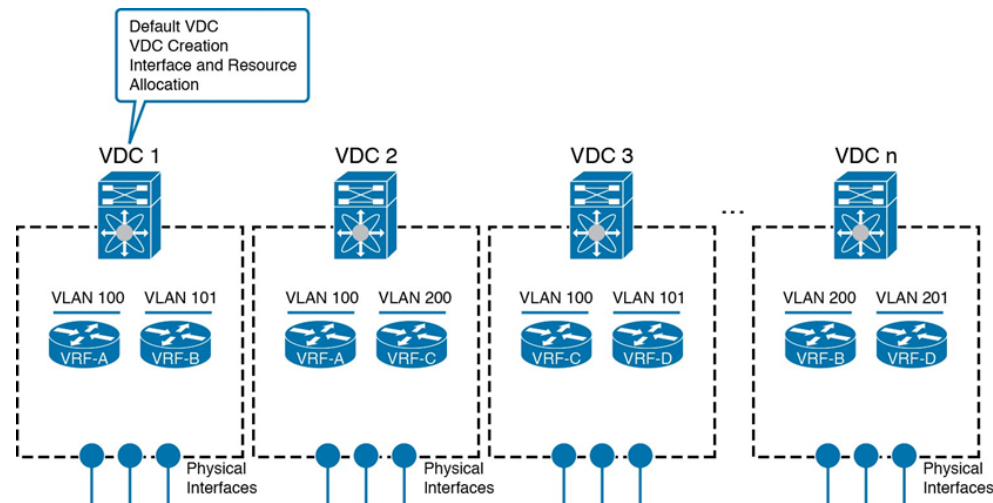


Source: Santana 2014

Student Questions

Virtual Device Context (VDC)

- ❑ Cisco proprietary concept.
- ❑ Combines concepts of virtual application delivery controllers (e.g., virtual firewalls) with Ethernet switching.
- ❑ A VDC includes VRFs and associated VLANs.
- ❑ A physical L3 switch can be partitioned into multiple VDCs



Ref: G. Santana, "Datacenter Virtualization Fundamentals," Cisco Press, 2014, ISBN: 1587143240

Washington University in St. Louis

<http://www.cse.wustl.edu/~jain/cse570-23/>

©2023 Raj Jain

Student Questions

- ❑ The L3 switch mentioned in the last bullet point means the device can do both L3 routing and L2 switching? *Yes.*
- ❑ Do physical devices that do VDCs have to be L3 switches? If so, is this why we can separate the switch to the L3 core router and the L2 aggregation switch? *Yes.*

VDC (Cont)

- ❑ A VDC has its own virtual data, control, and management plane.
 - Virtual Data Plane: Isolation between various VDCs
 - Virtual Control Plane: Own protocols, RIBs, ...
 - Virtual Management Plane: Managed by non-trusting entities. Different VDCs can have VRFs with the same name.
- ❑ Each VDC has its *physical interfaces*
A port can not belong to multiple VDCs.
⇒ Allows different VDCs to have overlapping VLAN IDs
- ❑ Large switches can have several hundred interfaces.
These switches can be partitioned using VDCs.


Student Questions

- ❑ What is the use of the management plane?
What is the difference between the management plane and the control plane?

Please wait till the SDN module.

VDC Example

- ❑ One switch can be used as an aggregation switch, and a core router
- ❑ Example only. It's not a common case.




Source: Santana 2014

Student Questions

- ❑ Is the VDC a router or a switch?
L3 Switch

Summary



- ❑ HSRP allows multiple routers to act as default gateway with the same VIP.
- ❑ VRRP is the IETF standard version of Cisco proprietary HSRP.
- ❑ VRF allows partitioning a router for multiple tenants.
- ❑ VDC allows partitioning a switch or application delivery controllers.

Student Questions

Reading List

- ❑ G. Santana, “Datacenter Virtualization Fundamentals,” Cisco Press, 2014, ISBN: 1587143240 (Safari Book)
- ❑ Javvin Technologies, "Network Protocols Handbook," Javvin Press, 2007, 380 pp., Chapter 5 and 55, ISBN:0974094528 (Safari book)

Student Questions

Wikipedia Links

- ❑ http://en.wikipedia.org/wiki/Hot_Standby_Router_Protocol
- ❑ http://en.wikipedia.org/wiki/Virtual_Router_Redundancy_Protocol
- ❑ http://en.wikipedia.org/wiki/Forwarding_information_base
- ❑ http://en.wikipedia.org/wiki/Routing_Information_Base
- ❑ http://en.wikipedia.org/wiki/Virtual_Routing_and_Forwarding

Student Questions

Acronyms

- ❑ BGP Border Gateway Protocol
- ❑ CE Customer edge
- ❑ FEX Fabric Extension
- ❑ FIB Forwarding Information Base
- ❑ GRE Generic Routing Encapsulation
- ❑ HSRP Hot Standby Router Protocol
- ❑ ID Identifier
- ❑ IETF Internet Engineering Task Force
- ❑ IP Internet Protocol
- ❑ MPLS Multi Protocol Label Switching
- ❑ OSPF Open Shortest Path First
- ❑ PE Provider Edge
- ❑ RFC Request for Comments
- ❑ RIB Routing Information Base
- ❑ RIP Routing Information Protocol
- ❑ VBE Virtual Bridge Port Extension

Student Questions

Acronyms (Cont)

- ❑ VDC Virtual Device Context
- ❑ VEB Virtual Edge Bridge
- ❑ VIP Virtual IP Address
- ❑ VLAN Virtual Local Area Network
- ❑ vPC Virtual Port Channels
- ❑ VRF Virtual Routing and Forwarding
- ❑ VRRP Virtual Router Redundancy Protocol
- ❑ VSS Virtual Switching System
- ❑ WAN Wide Area Network

Student Questions

Scan This to Download These Slides



Raj Jain

<http://rajjain.com>

Student Questions

http://www.cse.wustl.edu/~jain/cse570-23/m_09brv.htm

Related Modules



CSE567M: Computer Systems Analysis (Spring 2013),

https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),

https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcg5e_10TiDw



Wireless and Mobile Networking (Spring 2016),

https://www.youtube.com/playlist?list=PLjGG94etKypKeb0nzyN9tSs_HCd5c4wXF

CSE571S: Network Security (Fall 2011),

<https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u>



Video Podcasts of Prof. Raj Jain's Lectures,

<https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw>

Student Questions