

L102 - First Hop Redundancy

Network Design First Hop

First Hop Redundancy, Server Redundancy

Agenda

- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

L102 - First Hop Redundancy

First Hop Redundancy (Layer 3) **1**

- **The problem:**
 - How can local routers be recognized by IP hosts?
 - Note: Normally IP host has limited view of topology
 - IP host knows to which IP subnet connected
 - IP host knows one "Default Gateway" to reach other IP networks
 - Static configuration of "Default Gateway":
 - Loss of the default router results in a catastrophic event, isolating all end-hosts that are unable to detect any alternate path that may be available
- **Two design philosophies:**
 - Solve the problem at the IP host level
 - OS of the IP host need to support certain functionality in a appropriate way
 - Solve the problem at the IP router level
 - OS of the IP host need to support the basic functionality only
 - that is static configuration of one "Default Gateway"
 - Proprietary functionality may be needed at the router

© 2011, D.I. Manfred Lindner First Hop Redundancy, v1.7 3

First Hop Redundancy (Layer 3)

2

- **Methods for solving it at the IP host level:**
 - Proxy ARP
 - IDRP
 - DHCP
 - IP Routing (RIPv2, OSPF)
- **Methods for solving it at the IP router level:**
 - HSRP
 - VRRP
 - GLBP

L102 - First Hop Redundancy

Agenda

- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

5

IP Host, Default Gateway, ARP

- **IP routing from the view point of an IP host**
 - direct versus indirect delivery
 - depends on destination net-ID
 - net-ID equal to source net-ID -> direct delivery
 - » destination is in the same LAN
 - net-ID unequal to source net-ID -> indirect delivery
 - » destination is not in the same LAN
 - IP hosts are responsible for direct delivery of IP datagram's
 - IP hosts are responsible for choosing a default router ("default gateway") as next hop in case of indirect delivery of IP datagram's
 - **Transport of IP over a L2 technology**
 - L2 addresses (MAC addresses) must be used
 - mapping between IP and MAC address is done by Address Resolution Protocol (ARP)

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

6

L102 - First Hop Redundancy

Address Resolution Protocol

1

- **The mapping between MAC- and protocol-address on a LAN can be static (table entries) or dynamic (ARP protocol and ARP cache)**
- **Operation of ARP:**
 - Station A wants to send to station B and doesn't know the MAC address (both are connected to the same LAN)
 - A sends an ARP request in form of a MAC broadcast (destination = FF, source = Mac_A), ARP request holds IP address of B
 - Station B sees the ARP request with its IP address and sends an ARP reply as a MAC frame (SA=Mac_B, DA=Mac_A), B puts the newly learned mapping (source MAC- and IP-address of A) into its ARP cache

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

7

Address Resolution Protocol

2

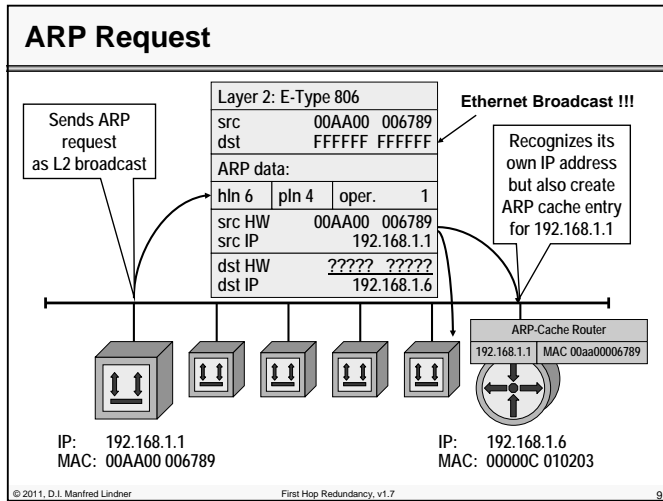
- The ARP reply holds MAC address of station B
- A stores the MAC- / IP-address mapping for station B in its ARP cache
- For subsequent packets from A to B or from B to A the MAC addresses are taken from the ARP cache (no further ARP request / reply)
- Entries in the ARP cache are deleted if they aren't used for a defined period (usually 5 min), this aging mechanism allows for changes in the network and saves table space
- ARP requests / reply are sent in Ethernet II (Type field 0x0806) or SNAP frames
- ARP has been designed to support different layer 3 protocols

© 2011, D.I. Manfred Lindner

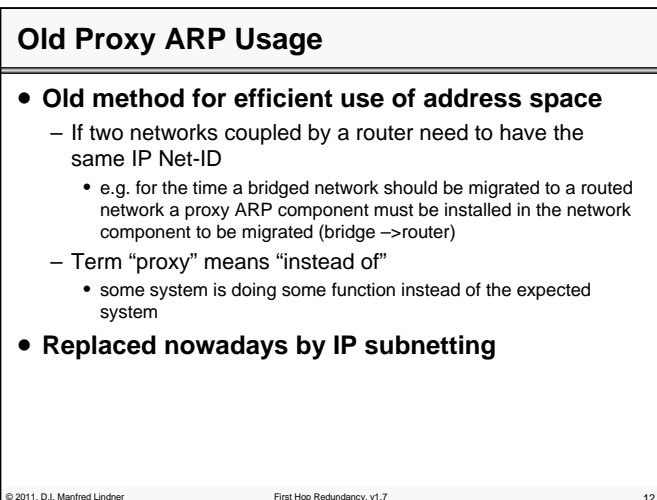
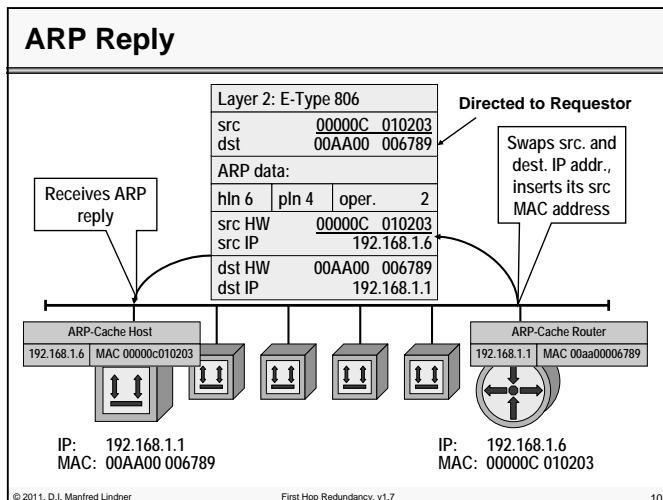
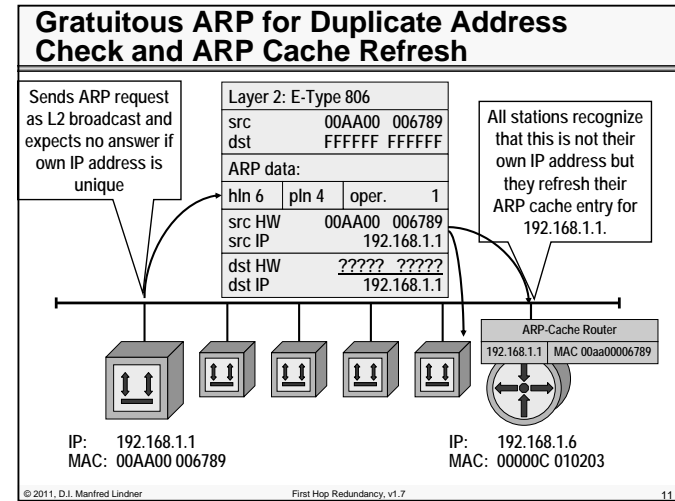
First Hop Redundancy, v1.7

8

L102 - First Hop Redundancy



L102 - First Hop Redundancy



L102 - First Hop Redundancy

Proxy ARP

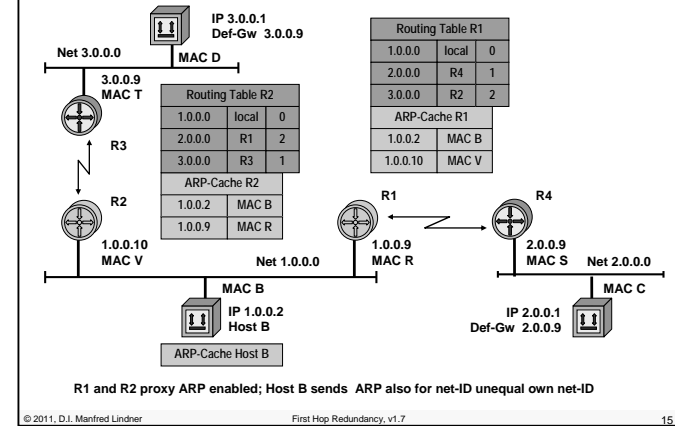
- In such as migration phase
 - Proxy ARP gives IP Hosts on both sides the impression of a homogenous network (single LAN)
- ARP requests for a station on the "other side"
 - are responded to on behalf of the destination with the MAC address of the Proxy system (the router in our case)
- IP datagram's will be sent to the layer 2 address (MAC) of the Proxy system
 - who forwards them to the destination network
- Important ARP rule:
 - Normal IP Protocol stack sends out an ARP Request for stations located on same subnet (same NET-ID) only !!!

Proxy ARP Usage Nowadays

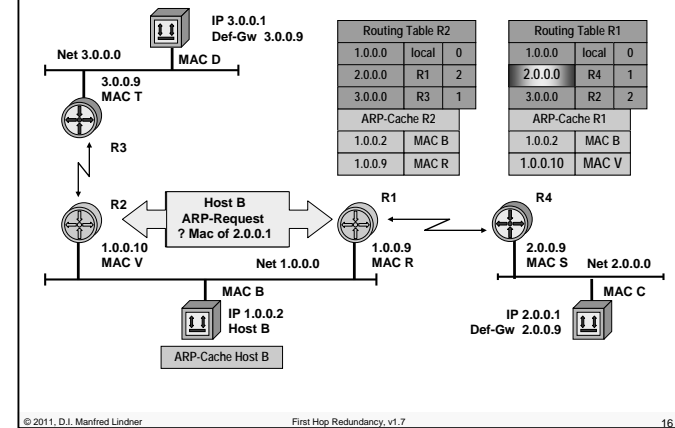
- Proxy ARP is can be used if an IP host didn't know the address of the default gateway or find it out dynamically:
 - In an IP host normally a static entry will tell the IP address of the router
 - if an IP datagram has to be sent to a non-local Net-ID, an ARP request will find the MAC address of the default gateway
 - With Proxy ARP extensions in the IP host and in the router
 - the MAC address of the router can be found without knowing the routers IP address
 - An ARP request will be sent for IP hosts with NET-IDs different from the local Net-ID and the router will respond
 - With Unix stations or Windows NT/XP:
 - proxy ARP extensions are triggered by setting the default gateway to the systems IP address itself

L102 - First Hop Redundancy

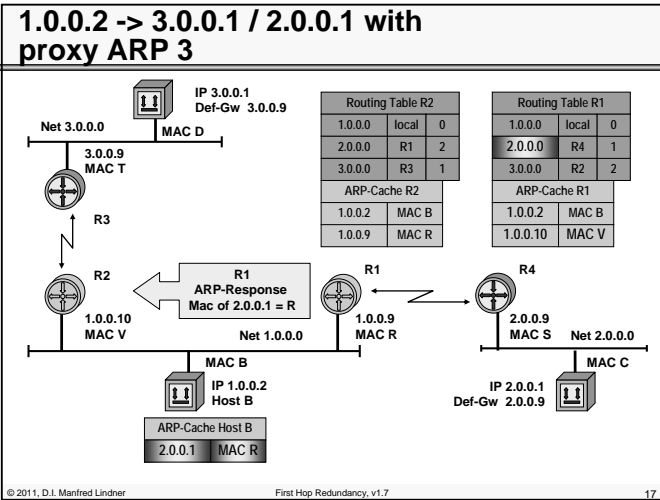
1.0.0.2 -> 3.0.0.1 / 2.0.0.1 with proxy ARP 1



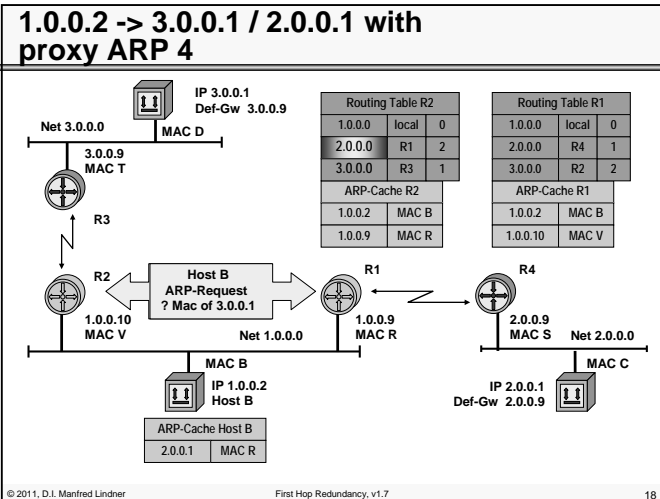
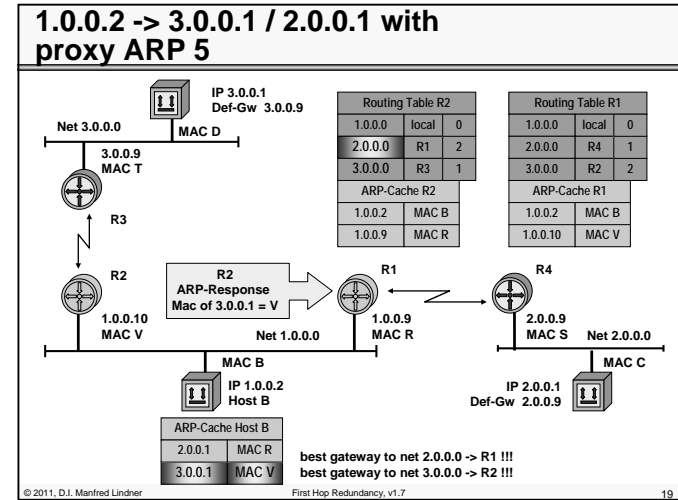
1.0.0.2 -> 3.0.0.1 / 2.0.0.1 with proxy ARP 2



L102 - First Hop Redundancy



L102 - First Hop Redundancy



Other Techniques to Solve the Problem 1

- IRDP
 - ICMP Router Discovery Messages (RFC 1256)
 - Routers periodically advertise their IP address on a shared media together with a preference value and a lifetime
 - ICMP Router Advertisement Message
 - Hosts may listen to these messages to find out all possible Default Gateways
 - or may ask by sending an ICMP Router Solicitation Message
- DHCP
 - Dynamic Host Configuration Protocol (RFC 2131)
 - More than one Default Gateway can be specified
 - Every Default Gateway has a preference value

L102 - First Hop Redundancy

Other Techniques to Solve the Problem 2

- **With IRDP and DHCP**

- You still depend on OS functionality in order to trigger switchover between redundant local routers
 - How often the currently selected router will be tested for reachability? What is if the currently selected router is reachable via LAN but networks behind are not reachable?

- **Therefore running a classical IP routing protocol on the IP host would be optimal**

- RIPv2
 - But slow convergence if the currently selected router fails, no hello messages hence 180 seconds for recognizing that event
- OSPF
 - Fast convergence because of hello messages, the best but the most complex solution

Agenda

- **First Hop Redundancy**

- Proxy ARP, IDRP, DHCP
- HSRP
- VRRP
- GLBP
- Design Access WAN

- **Server Load Balancing**

- SLB
- DNS

L102 - First Hop Redundancy

HSRP – Hot Standby Router Protocol

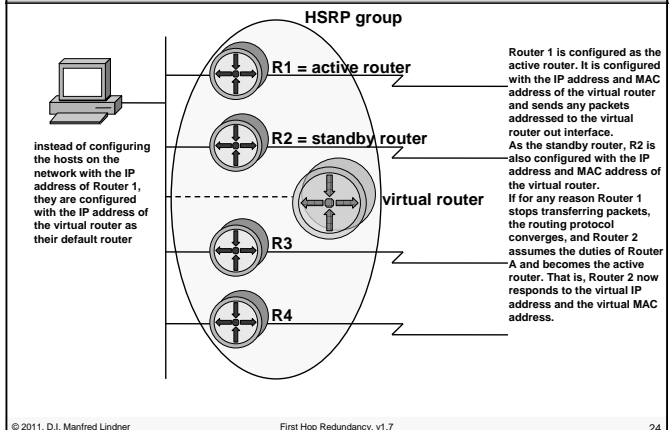
- **HSRP (Hot Standby Router Protocol)**

- Proprietary protocol invented by Cisco
- RFC 2281 (Informational)

- **Basic idea: a set of routers present the illusion of a single virtual router to the hosts on the LAN**

- Active router
 - one router is responsible for forwarding the packets that hosts send to the virtual router
- Standby router
 - if active router fails, the standby assumes the packet forwarding duties of the active router
- Conspiring routers form a HSRP group

Terminology



L102 - First Hop Redundancy

HSRP Operation

1

- **Principle:**

- A group of routers forms a HSRP group
- The group is represented by a virtual router
 - With a virtual IP address and virtual MAC address for that group
- IP hosts are configured with the virtual IP address as default gateway
- One router is elected as the active router, one router is elected as the standby router of that group
- Active router responds to ARP request directed to the virtual IP address with the virtual MAC address
- Standby router supervise if the active router is alive and can take over the role of the active router
 - HSRP protocol using UDP messages to port 1985, IP multicast 224.0.0.2, and Ethernet multicast as destination address (HSRP version 1)
- Router must be able to support more than one unicast MAC address on an Ethernet interface

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

25

HSRP Operation

2

- **Roles of routers:**

- **Standby router**
 - The backup router in case the active router fails for the subnet
 - In that case, the standby router becomes the active router and starts forwarding traffic destined to the virtual IP address
- **Active router**
 - The active router forwards traffic destined to the virtual IP address
- **Additional HSRP member routers - Other**
 - Other routers are neither active nor standby; they just monitor the messages of the current active and standby routers and transition into one of those roles if the current router fails for the subnet
- **Virtual router**
 - The virtual router is not an actual router
 - Rather, it is a concept of the entire HSRP group acting as one virtual router as far as hosts on the subnet are concerned

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

26

L102 - First Hop Redundancy

HSRP Operation

3

- **Roles of router (cont.):**

- Active, Standby, Other are defined by HSRP priority
- HSRP Priority value can be configured
 - Default value is 100
- The higher the better
 - Will become the active router after initialization
 - If priority is equal than the higher IP address decides (tiebreaker)
- Preempt allows to give up the role of the active router when a router with higher priority is activated or reported
 - e.g. a failed router comes back or tracking has changed priority

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

27

HSRP Operation

4

- **Failover scenarios:**

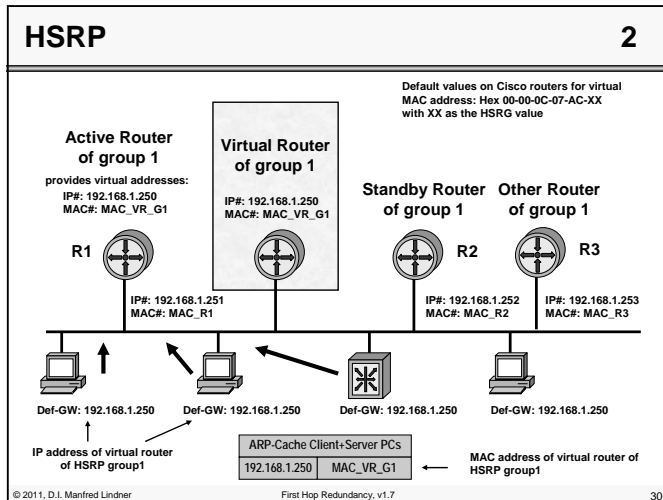
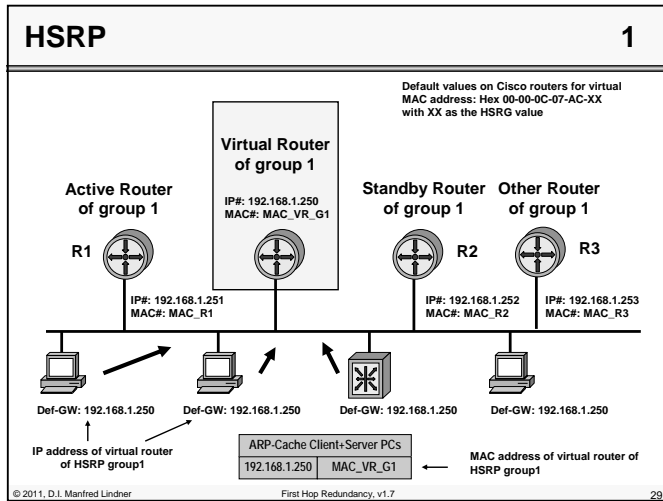
- Active router not reachable via LAN
 - Standby router will take over active role
 - A new standby router is elected from the remaining routers of a HSRP group
 - Timing depends on HSRP hello message interval and hold-time
 - Default hello-time = 3 seconds, default hold-time = 10 seconds
- Active router losses connectivity to a WAN interface (basic tracking options) or losses connectivity to an IP route (enhanced tracking options)
 - If tracking and preempt is configured standby router will take over
 - Tracking will lower the priority
 - Preempt allows another router to take over the role of the active router even if the current active router does not fail
- Enhanced tracking options depend on IOS version

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

28

L102 - First Hop Redundancy



L102 - First Hop Redundancy

- HSRP**
3

 - The **active router** assumes and maintains its active role through the transmission of hello messages (default 3 seconds, HSRP version 1)
 - The hello interval time defines the interval between successive HSRP hello messages sent by active and standby routers
 - The router with the highest standby priority in the group becomes the active router
 - The default priority for an HSRP router is 100
 - When the preempt option is not configured, the first router to initialize HSRP becomes the active router

© 2011, D.I. Manfred Lindner First Hop Redundancy, v1.7 31

- HSRP**
4

 - The second router in the HSRP group to initialize or second highest priority is elected as the **standby router**
 - The function of the standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router becomes inoperable
 - The standby router also transmits hello messages to inform all other routers in the group of its standby router role and status


© 2011, D.I. Manfred Lindner First Hop Redundancy, v1.7 32

L102 - First Hop Redundancy

HSRP	5
<ul style="list-style-type: none"> • The virtual router presents a consistent available router (default gateway) to the hosts • The virtual router is assigned its own IP address and virtual MAC address <ul style="list-style-type: none"> – however, the active router acting as the virtual router actually forwards the packets • Additional HSRP member routers - other routers : <ul style="list-style-type: none"> – These routers in listen state monitor the hello messages but do not respond – They forward any packets addressed to their own IP addresses. – They do <u>not</u> forward packets destined for the virtual router because they are not the active router. 	
© 2011, D.I. Manfred Lindner	First Hop Redundancy, v1.7 33

L102 - First Hop Redundancy

HSRP	7
<ul style="list-style-type: none"> • When the active router fails, the HSRP routers stop receiving hello messages from the active and the standby router assumes the role of the active router <ul style="list-style-type: none"> – This occurs when the <i>holdtime</i> expires (default 10 seconds, HSRP version 1) – If there are other routers participating in the group, those routers then contend to be the new standby router • Because the new active router assumes both the IP address and virtual MAC address of the virtual router, the end stations see no disruption in service <ul style="list-style-type: none"> – The end-user stations continue to send packets to the virtual router's virtual MAC address and IP address where the new active router delivers the packets to the destination 	
© 2011, D.I. Manfred Lindner	First Hop Redundancy, v1.7 35



HSRP Protocol Fields				
<ul style="list-style-type: none"> • standby protocol runs on top of UDP (port 1985) <ul style="list-style-type: none"> – IP packets are sent to multicast address 224.0.0.2 or 224.0.0.102 with a IP TTL = 1 				
0	4	8	16	31
Version	Op Code	State	HelloTime	
Holdtime	Priority	Group	Reserved	
Authentication Data				
Authentication Data				
Virtual IP Address				
<ul style="list-style-type: none"> - version: version of the HSRP messages - op code: 3 types <ul style="list-style-type: none"> hello: indicates that a router is running and is capable of becoming the active or standby router coup: when a router wishes to become the active router resign: when a router no longer wishes to be the active router 		<ul style="list-style-type: none"> - states: initial, learn, listen, speak, standby, active - hellotime: contains the period between the hello messages that the router sends - holdtime: amount of time the current hello message is valid - priority: compares priorities of 2 different routers - group: identifies standby group (0..255) - authentication data: cleartext 8 character reused password 		
© 2011, D.I. Manfred Lindner	First Hop Redundancy, v1.7	36		

L102 - First Hop Redundancy

HSRP Versions

- **HSRP version 1:**
 - Second timers
 - 256 groups (0 – 255)
 - Virtual Mac Address: 00-00-0C-07-AC-XX
 - XX value = group number
 - IP multicast 224.0.0.2
- **HSRP version 2:**
 - Millisecond timers
 - Hello-time 15 - 999 msec
 - Hold-time - 3000 msec
 - 4096 groups (0-4095)
 - Allow a group number to match the VLAN-ID
 - Virtual Mac Address: 00-00-0C-9F-FX-XX
 - X-XX value = group number
 - IP multicast 224.0.0.102
 - To avoid conflicts with CGMP (Cisco Group Management Protocol) by using 224.0.0.2

HSRP States Details

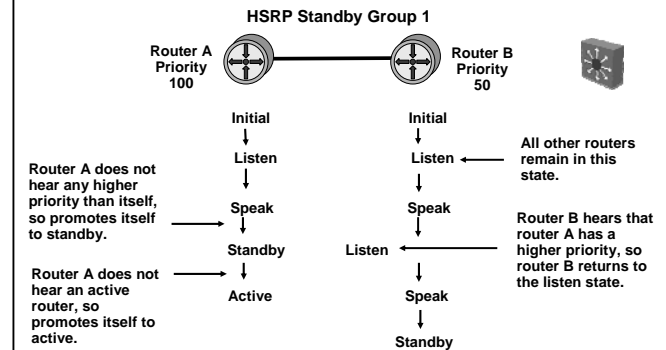
1

- **Initial state**— All routers begin in the initial state. This state is entered via a configuration change or when an interface is initiated.
- **Learn state**— The router has not determined the virtual IP address, and has not yet seen a hello message from the active router. In this state, the router is still waiting to hear from the active router.
- **Listen state**— The router knows the virtual IP address, but is neither the active router nor the standby router. All other routers participating in the HSRP group besides the active or standby routers reside in this state.
- **Speak state**— HSRP routers in the speak state send periodic hello messages and actively participate in the election of the active or standby router. The router remains in the speak state unless it becomes an active or standby router.
- **Standby state**— In the standby state, the HSRP router is a candidate to become the next active router and sends periodic hello messages. There must be at least one standby router in the HSRP group.
- **Active state**— In the active state, the router is currently forwarding packets that are sent to the virtual MAC and IP address of the HSRP group. The active router also sends periodic hello messages.

L102 - First Hop Redundancy

HSRP States Details

2

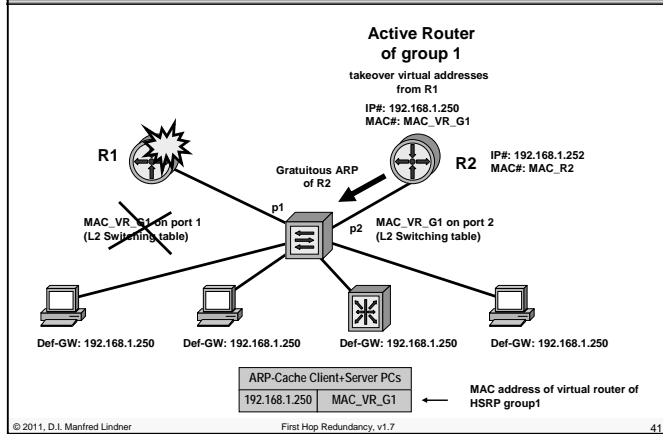


HSRP Additional Aspects

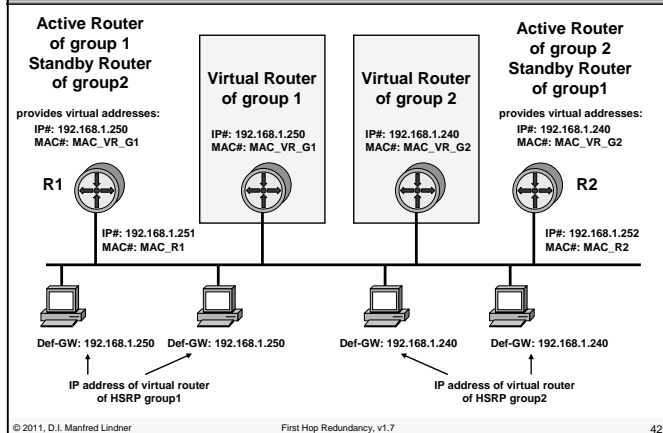
- **L2 Switching Table Refresh:**
 - Done by gratuitous ARP in case of switchover
- **Load Balancing:**
 - You can achieve this by specifying at least two different HSRP groups with complementary roles
- **HSRP Security**
 - Authentication of messages by generation of fingerprints and checking this fingerprints
 - Based on keyed MD5
 - Against HSRP spoofing

L102 - First Hop Redundancy

HSRP – Gratuitous ARP



HSRP Load Balancing



L102 - First Hop Redundancy

Agenda

- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

VRRP Operation

1

- **VRRP (Virtual Router Redundancy Protocol)**
 - RFC 2338 (Standards Track)
- **Principle:**
 - A group of routers forms a VRRP group
 - The group is represented by a virtual router
 - With is identified by a VRID (Virtual Router ID) and a virtual MAC address
 - One router is elected as the **virtual router master**, all other routers get the role of **virtual router backup** routers
 - The real IP address of the virtual router master become the IP address of the virtual router for a given VRRP group
 - IP address owner
 - Default Gateway of IP hosts is configured with the IP address of the virtual router for a given VRRP group
 - Virtual router master responds to ARP request directed to the IP address of the virtual router with the virtual MAC address
 - Backup routers supervise if master router is alive and take over the role of the master in case of failure
 - VRRP protocol using IP protocol number 112, IP multicast 224.0.0.18, and Ethernet multicast as destination address
 - Router must be able to support more than one unicast MAC address on an Ethernet interface

L102 - First Hop Redundancy

VRRP Operation

2

• **Roles of router:**

- Virtual router master, virtual router backup defined by VRRP priority
 - Priority value can be configured
 - Default value is 100
 - The higher the better
 - Will become the master after initialization
 - If priority is equal than the higher IP address decides
- Preempt allows to give up the role of the master router when a router with higher priority is activated or reported
 - e.g. a failed router comes back or tracking has changed priority

• **Load Balancing:**

- Specify at least two different VRRP groups with complementary roles

• **VRRP authentication:**

- Based on keyed MD5
- Against VRRP spoofing

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

45

VRRP Operation

3

• **Failover scenarios:**

- Master router not reachable via LAN
 - Backup router with highest priority will take over master role
 - Timing depends on VRRP advertisements interval and master down interval
 - Default advert-interval = 1 seconds
 - Default master-down-interval = 3 * advert-interval + skew-time
- Master router losses connectivity to a WAN interface (basic tracking options) or losses connectivity to an IP route (enhanced tracking options)
 - If tracking and preempt is configured backup router will take over
 - Tracking will lower the priority
 - Preempt allows another router to take over the role of the master router even if the current master router does not fail
- Enhanced tracking options depend on IOS version

© 2011, D.I. Manfred Lindner

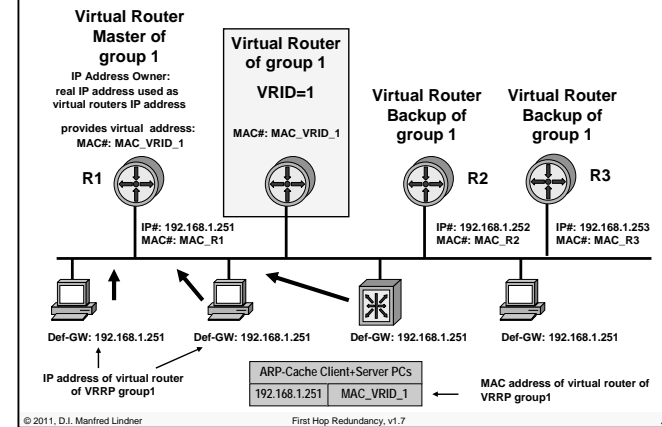
First Hop Redundancy, v1.7

46

L102 - First Hop Redundancy

VRRP

1



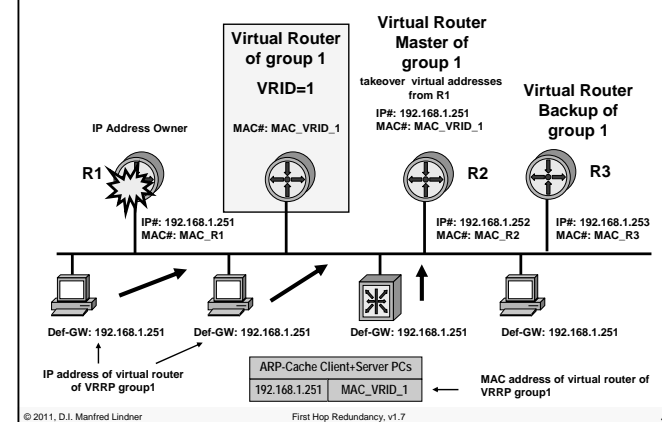
© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

47

VRRP

2

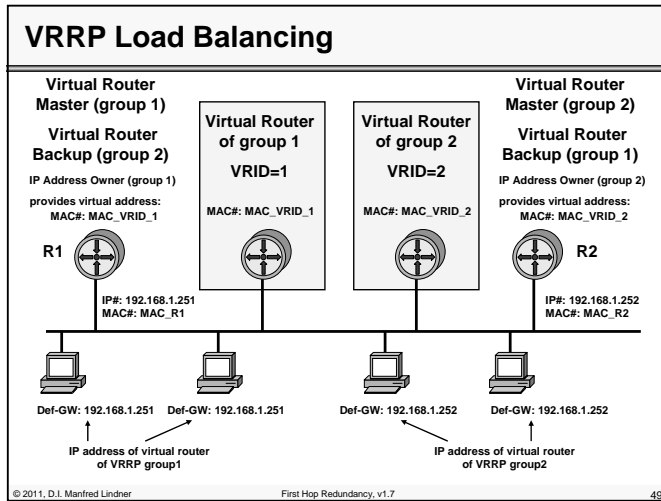


© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

48

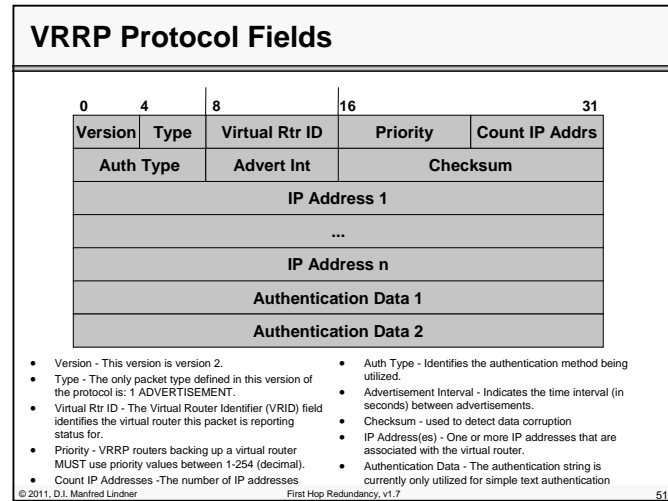
L102 - First Hop Redundancy



Some VRRP Details

- **VRRP:**
 - Second or milliseconds timers
 - VRID range
 - 1 – 255
 - Maximum 255 groups
 - Virtual Mac Address: 00-00-5E-00-01-VRID
 - VRID value = group number
 - IP multicast 224.0.0.18

L102 - First Hop Redundancy



Agenda

- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

L102 - First Hop Redundancy

GLBP Operation

1

- **GLBP (Gateway Load Balancing Protocol)**
 - Proprietary protocol invented by Cisco
 - Protects data traffic from a failed router or circuit, like Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP), while allowing packet load sharing between a group of redundant routers
- **Principle:**
 - A group of routers forms a GLBP group
 - The group is represented by a virtual router
 - With one virtual IP address and several virtual MAC addresses for that group
 - Members of a GLBP group elect one gateway to be the **active virtual gateway (AVG)** for that group
 - The other group members provide backup for the AVG in the event that the AVG becomes unavailable
 - The function of the AVG is to assign a unique virtual MAC address to each member of the GLBP group

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

53

GLBP Operation

2

- Each member of the GLBP groups assumes responsibility for forwarding packets sent to the virtual MAC address assigned to it by the AVG
- Hence these gateways are known as **active virtual forwarders (AVF's)** for their corresponding virtual MAC address
- Default Gateway of IP hosts is configured with the IP address of the virtual router for a given GLBP group
- The AVG is responsible for answering Address Resolution Protocol (ARP) requests for the virtual IP address
 - Load sharing is achieved by replying to the ARP requests with different virtual MAC addresses (different AVF's)
- Communication
 - GLBP protocol using UDP messages to port 3222, IP multicast 224.0.0.102, and Ethernet multicast as destination address
- Router must be able to support more than one unicast MAC address on an Ethernet interface

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

54

L102 - First Hop Redundancy

GLBP Operation

2

- **Roles of router:**
 - AVG, AVF defined by GLBP priority
 - Priority value can be configured
 - Default value is 100
 - The higher the better
 - Will become the AVG after initialization
 - If priority is equal than the higher IP address decides
 - Preempt allows to give up the role of the AVG when a router with higher priority is activated or reported
 - e.g. a failed router comes back or tracking has changed priority
 - Forwarder preempt allows to give up the role of the AVF when a router with higher priority is activated or reported

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

55

GLBP Operation

3

- **Load Balancing:**
 - Round-robin
 - ARP replies specify all AVF's in turn
 - Weighted
 - ARP replies specify AVF's based on the weight assigned to them
 - Host-dependent
 - ARP replies specify always the same AVF for a given host
- **GLBP authentication:**
 - Based on keyed MD5
 - Against GLBP spoofing

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

56

L102 - First Hop Redundancy

GLBP Operation

4

• Failover scenarios:

- AVG/AVF router not reachable via LAN
 - Other router will take over AVG/AVF role
 - Timing depends on GLBP hello message interval and hold-time
 - Default hello-time = 3 seconds, default hold-time = 10 seconds
- AVF router losses connectivity to a WAN interface (basic tracking options) or losses connectivity to an IP route (enhanced tracking options)
 - If tracking and preempt is configured standby router will take over
 - Tracking will lower the priority
 - Preempt allows another router to take over the role of the active router even if the current active router does not fail
- Enhanced tracking options depend on IOS version

© 2011, D.I. Manfred Lindner

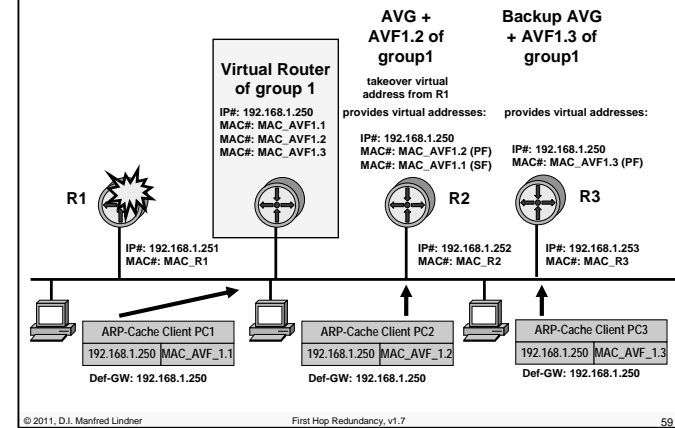
First Hop Redundancy, v1.7

57

L102 - First Hop Redundancy

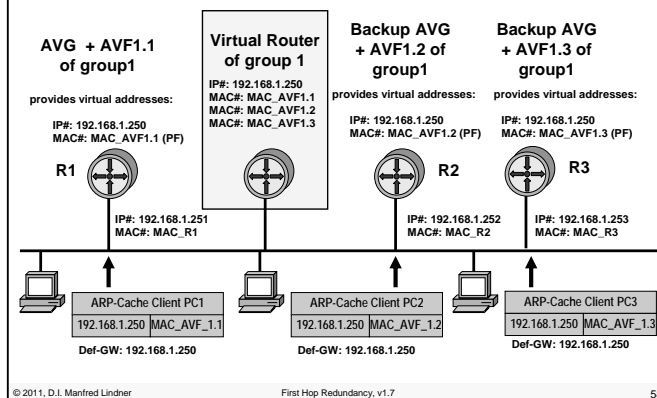
GLBP AVG-Failure

2



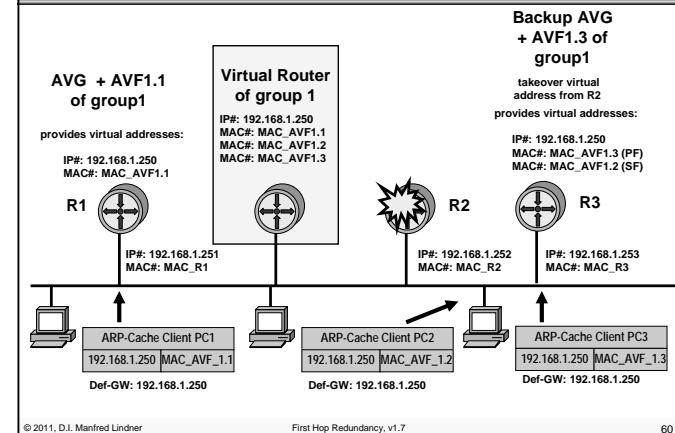
GLBP

1



GLBP AVF-Failure

3



L102 - First Hop Redundancy

GLBP Details

1

- **Second or milliseconds timers**
- **GLBP will use the following multicast destination for packets sent to all GLBP group members:**
 - 224.0.0.102, UDP port 3222
- **Protocol allows for 1024 groups and 255 forwarders**
 - Four virtual MAC addresses per group
- **Hardware restrictions limit actual number of groups and forwarders**
- **1024 groups (0 – 1023)**
- **Virtual MAC addresses will be of the form:**
 - 0007.b4yy.yyyy
 - where yy.yyyy equals the lower 24 bits; these bits consist of 6 zero bits, 10 bits that correspond to the GLBP group number, and 8 bits that correspond to the virtual forwarder number
 - 0007.b400.0102 : last 24 bits = 0000 0000 0000 0001 0000 0010 = GLBP group 1, forwarder 2

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

61

GLBP Details

2

- **Hello messages are exchanged between group members**
 - AVG election by priority
 - Virtual MAC distribution, learning of VF instances
- **If virtual MAC addresses was assigned by AVG**
 - Primary virtual forwarder (PF)
- **If virtual MAC addresses learned from hello messages**
 - Secondary virtual forwarder (SF)
- **Procedure in case of an AVF failure:**
 - GLBP migrates hosts away from the old forwarder number (PF-number) using two timers seen in hello messages sent by the AVG
 - Timers start as an SF takeover responsibility for a failed PF
 - Redirect timer, PF-unavailability timer
 - During the redirect time the AVG continues to redirect hosts to the old virtual forwarder MAC address; when the redirect time expires, the AVG stops using the old virtual forwarder MAC address in ARP replies
 - When the PF-unavailability timer fires, the virtual forwarder is removed from all gateways in the GLBP group; the expired virtual forwarder number becomes eligible for reassignment by the AVG; timer must be long enough to allow aging of ARP caches

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

62

L102 - First Hop Redundancy

GLBP Benefits

- **Load sharing**
 - traffic from LAN clients can be shared by multiple routers, thereby sharing the traffic load more equitably among available routers
- **Multiple virtual routers**
 - GLBP supports up to 1024 virtual routers (GLBP groups) on each physical interface of a router and up to four virtual forwarders per group
- **Authentication**
 - A router within a GLBP group with a different authentication string than other routers will be ignored by other group members

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

63

Agenda

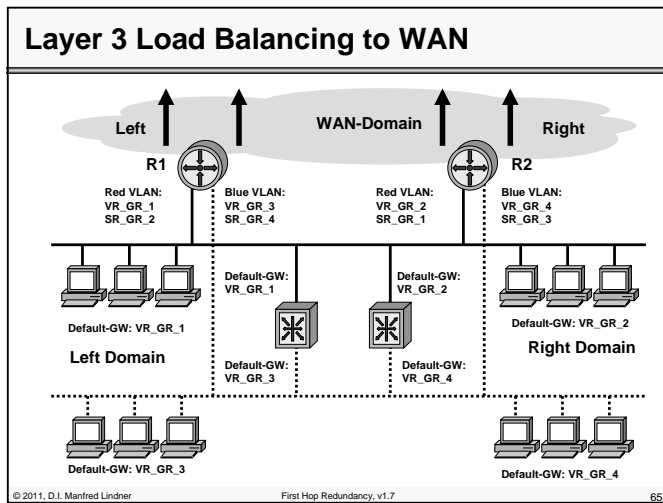
- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

© 2011, D.I. Manfred Lindner

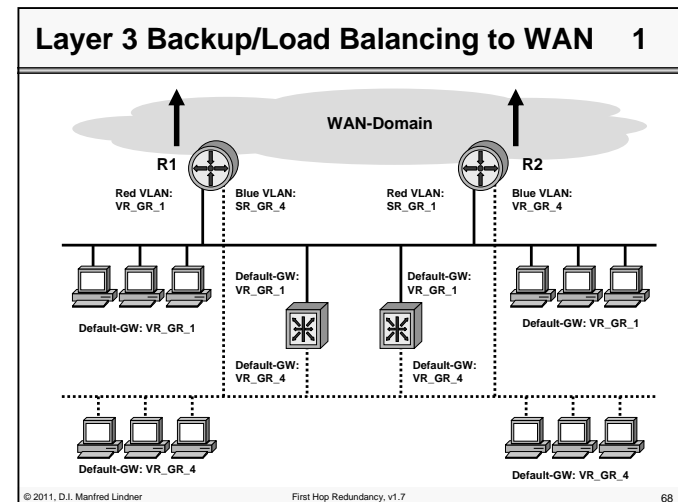
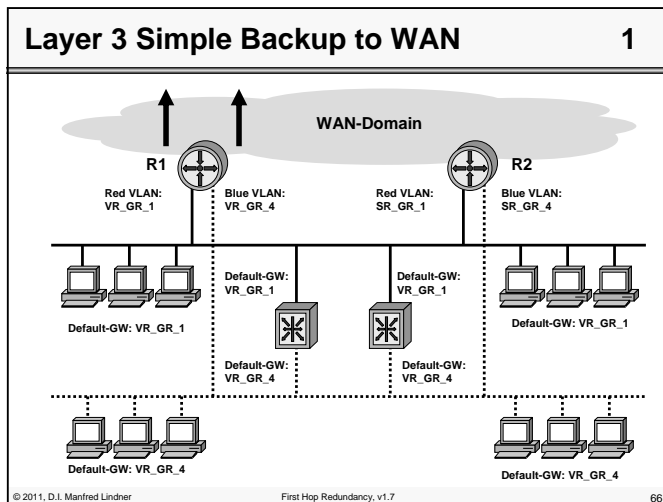
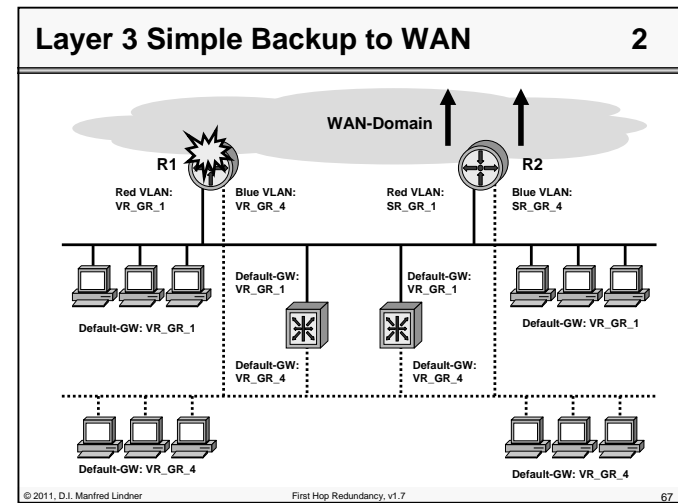
First Hop Redundancy, v1.7

64

L102 - First Hop Redundancy

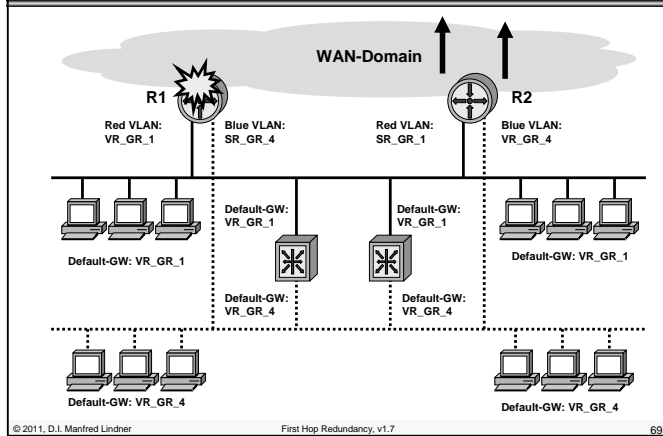


L102 - First Hop Redundancy



L102 - First Hop Redundancy

Layer 3 Backup/Load Balancing to WAN 2



Agenda

- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

L102 - First Hop Redundancy

SLB Operation 1

- **SLB (Server Load Balancing)**
 - Proprietary technology in Cisco IOS
- **Principle:**
 - A router represents a virtual server IP address and a corresponding TCP/UDP port handled by such a server
 - Clients work with this virtual server IP address
 - TCP connections, UDP requests
 - Client-to-server traffic is intercepted by the router and directed to real servers based on a scheduling algorithm
 - Destination address is changed (NAT)
 - Server-to-client traffic is intercepted by the router to change the source address of the packet to the virtual server IP address (NAT)

SLB Operation 2

- Load balancing can be
 - Weighted round-robin
 - A real server get TCP connections based on the weight
 - Weight = 1 means real round-robin
 - Weighted least connections
 - Decision is based on the current active connections to a server
 - The server with the least active connections get the next new request
 - Weight will influence the decision in such a way, that a server with a larger weight can take more sessions than a server with a smaller number
 - Weight = 1 means simple least connection
 - Client-assigned load balancing
 - You can specify which client IP addresses / IP subnets are permitted to use a virtual server
 - Hence different IP subnets can be directed to different server farms

L102 - First Hop Redundancy

SLB Operation

3

- Sticky based connections
 - For a given IP client address a new connection from this client is directed to the same real server as former connections
 - Sticky timer to control how long the history state remains in the router
 - Important for HTTP (port 80)
 - Important if HTTP (port 80) and HTTPS (port 443) are coupled
- Automatic Failover Detection
 - If a TCP connection attempt to real server fails a failure counter is incremented and when a configurable threshold is reached the server is removed from the active server list
- Automatic Unfail
 - A server removed from the active list is tried with a new connection after timeout of a retry timer; if successful then placed back on the list of active servers otherwise retry timer is restarted and server remains out of service
- Dynamic Feedback Protocol
 - Host agent report their performance to the router / SLB in form of a weight; reported weight is used as input for scheduling decision

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

73

SLB Operation

4

- SynGuard
 - Limits rate of TCP SYN segments to the virtual server to prevent SYN Flood Denial-of-service attacks
- NAT
 - Directed Mode
 - Virtual server IP address need to be translated by NAT
 - Performance impact on the router / SLB (L3 and L4 modified !!!)
 - Must be used if there is no direct L2 connectivity between the router and the real servers
 - Dispatched mode
 - Virtual server IP address is configured as loopback / secondary address on each of the real server, hence client packets are accepted by the server and server packets can use the loopback address as source
 - In such a case no NAT need to be performed at the router / SLB hence better performance
 - Requirement is a direct L2 connectivity between the router and the real server
 - Works because packets are addressed via L2 MAC addresses to reach the real server

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

74

L102 - First Hop Redundancy

SLB Operation

5


- IOS SLB Stateless Backup
 - IOS SLB can work together with HSRP to avoid single point of failures
 - Backup of the virtual server towards the clients

© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

75

SLB Basic Function

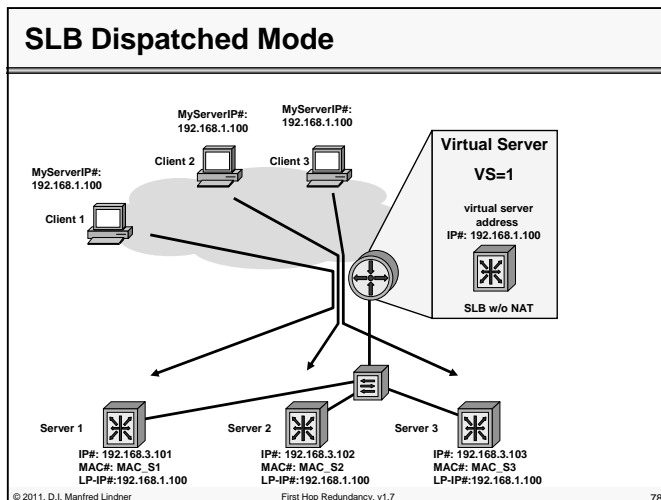
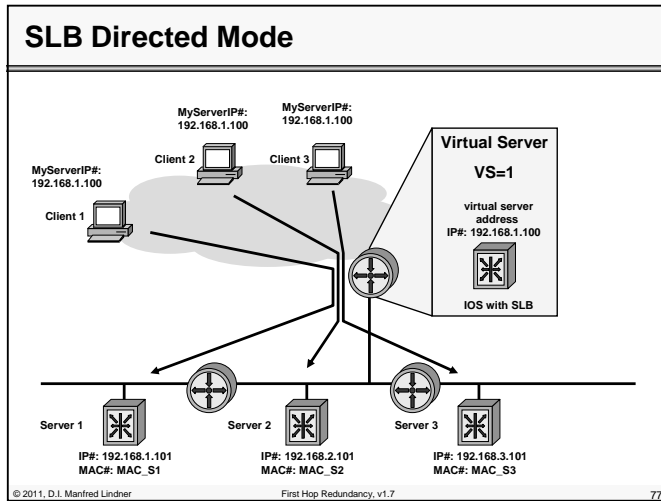


© 2011, D.I. Manfred Lindner

First Hop Redundancy, v1.7

76

L102 - First Hop Redundancy



L102 - First Hop Redundancy

Agenda

- **First Hop Redundancy**
 - Proxy ARP, IDRP, DHCP
 - HSRP
 - VRRP
 - GLBP
 - Design Access WAN
- **Server Load Balancing**
 - SLB
 - DNS

SLB with DNS

- **Server Load Balancing by DNS**
 - Clients asking for the same server name will be provided with different IP addresses
 - Often for free in available DNS SW package
 - Drawbacks
 - How does a DNS server detect a server failure?
 - Dynamic DNS?
 - How long will be the DNS TTL for the resolved name?
 - DNS cache
 - Do not mistake it with HTTP redirection
 - HTTP Get request redirected after inverse DNS lookup to a local server of a company
 - Inverse DNS lookup finds out the domain to a given client IP address
 - If .at -> assumes that user wants access to company.at server