

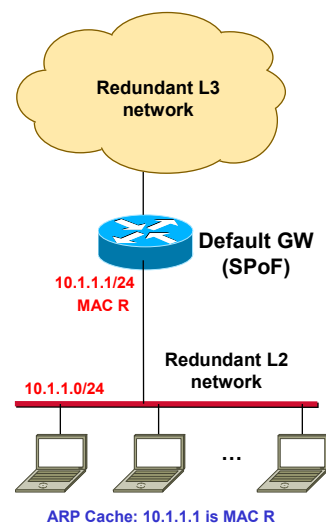
## First hop redundancy

HSRP, VRRP, CARP, GLBP, Proxy ARP  
Route Tracking and SLA Measurements

## The first-hop problem



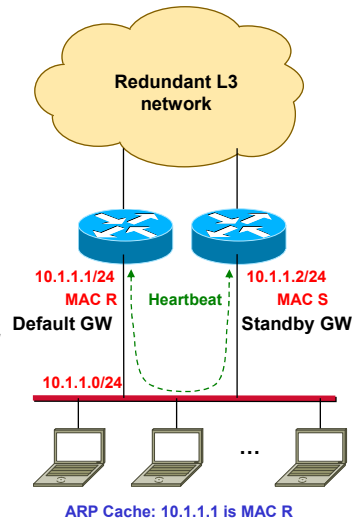
- LAN consists of a number of transparent bridges
  - ◆ Bridges have redundant interconnections
  - ◆ Spanning Tree ensures that end systems have valid paths to each other
- At least one Default Gateway per LAN (typically: per VLAN)
- But this one is single point of failure (SPoF)



## The first-hop problem (cont.)




- We want multiple routers to represent the single "Default Gateway"
  - ♦ Note that the actual delivery (i. e. the traffic path) is realized via MAC addresses
- Therefore we require that the IP and MAC addresses of the Default Gateway remains constant
  - ♦ Otherwise existing sessions would be interrupted
  - ♦ All existing ARP cache entries in the end systems must remain valid
- Redundant Gateways must know of each other
  - ♦ Check mutual liveliness
  - ♦ Agreement upon Default-GW role




## Typical Procedure



### Before Failover



### After Failover



## Path consistency



- **Only bridging tables must be (quickly) updated**
  - ♦ Automatically done via heartbeat packets sent by actual Default Gateway
  - ♦ Therefore, high heartbeat frequency desired (ms)
  - ♦ Ideally no packet lost
- **Neighbour routers must also learn which one of the Failover-Pair provides connectivity to the LAN**
  - ♦ In some cases asymmetric paths would be acceptable (but requires availability of standby router)
  - ♦ In case of a node failure, symmetric paths are mandatory!
  - ♦ Announcements either via routing protocol or also via transparent failover solution

## Failure Types



- **Link failure**
  - ♦ Detectable via mutual heartbeat
- **Node failure**
  - ♦ Detectable via mutual heartbeat
  - ♦ Requires symmetric paths
- **Route failure**
  - ♦ Default Gateway cannot forward packets from/to WAN
  - ♦ Either because of WAN interface failure (solution: interface tracking)
  - ♦ Or path is broken failure (solution: route tracking)

# Hot Standby Routing Protocol (HSRP)

(C) Herbert Haas 2005/03/11

## Overview

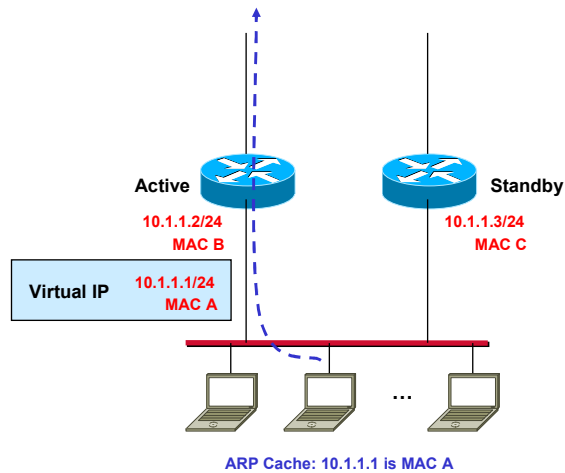


- Invented by Cisco, RFC 2281
- Several routers may be member of a HSRP group
- The whole HSRP group resembles a virtual router
  - ♦ Represented by a single virtual IP and a single virtual MAC address
- Three roles: Active, Standby, Other
  - ♦ Exactly one router is active only
  - ♦ Exactly one router is standby only
- The active/standby roles are chosen by a configurable priority value (default=100, higher is better, 0..255)
  - ♦ highest IP address is used as tie breaker

(C) Herbert Haas 2005/03/11

8

## HSRP Example



## Procedures



- Only the active router accepts packets destined for the virtual IP and MAC address
  - ♦ Obviously this router responds to ARP requests for the virtual IP address
- If the standby router is no longer available (because it either failed or has become active) then another router becomes standby
- Failure detection times
  - ♦ Version 1: Default hello time: 3 sec, Default hold time: 10 sec
  - ♦ Version 2: Hello time: **15-999 msec**, Hold time: up to 3000 msec

## Note



- **All non-active routers are still reachable and theoretically usable**
  - ♦ Ping works
  - ♦ Static routes possible
- **ICMP Redirect**
  - ♦ Must be disabled because it may reveal a routers primary address
- **If Proxy ARP is needed**
  - ♦ All HSRP routers must response with the virtual MAC address – regardless of their HSRP state!
  - ♦ Reason: Suppressing a Proxy ARP response could result in lack of any proxy ARP response being generated

## Optional Preemption



- **When a router with higher priority is added to the HSRP group it will immediately become active**

## Security



- **Any router redundancy protocol should use message authentication!**
  - ◆ Otherwise DoS from inside possible
  - ◆ DoS from outside rather impossible since only link-local multicast addresses used
- **Standard RFC 2281 only supports plaintext authentication**
  - ◆ Should only prevent misconfigurations
- **Since Cisco IOS 12.3(2)T also MD5 authentication supported**

## Failover Scenarios



- **Active router not reachable after hold time**
- **If *Object Tracking* is enabled, the priority can be reduced based on "lost properties", e. g.**
  - ◆ **Active router loses connectivity of one WAN interface**
    - "Interface tracking" – Only routers with active WAN interface (or other outside interface) may be potential default gateway
  - ◆ **Active router loses connectivity to one IP route**

```
! Example: When the IP routing capability of serial0/1 fails, the  
! HSRP priority is decremented by 10 so that the standby router  
! becomes active  
track 100 interface serial0/1 ip routing  
!  
interface Ethernet0/0  
 ip address 10.1.1.2 255.255.255.0  
 standby 1 preempt  
 standby 1 ip 10.1.1.1  
 standby 1 priority 105  
 standby 1 track 100 decrement 10
```

## Load Balancing



- **Specify at least two different HSRP groups**
  - ◆ HSRP version 1 supports up to 256 groups (0..255)
  - ◆ HSRP version 2 supports up to 4096 groups (0..4095)
- **One for each VLAN**

## Protocol details



- **Uses UDP port 1985**
- **Uses IP multicast addresses**
  - ◆ Version 1: 224.0.0.2 (all local routers)
  - ◆ Version 2: 224.0.0.102 (to avoid conflict with CGMP)
- **Optionally keyed MD5 Authentication**
  - ◆ Relatively weak!
  - ◆ Unprotected HSRP is a critical security hole (DoS prone)
- **Cisco's default virtual MAC address:**
  - ◆ Version 1: 00-00-0C-07-AC-XX
  - ◆ Version 2: 00-00-0C-9F-FX-XX
  - ◆ The XX identifies the group number

## Message Types



- **Hello**
  - ◆ This is the keep-alive packet
  - ◆ Contains priority number
- **Coup**
  - ◆ Sending router wishes to become active router
- **Resign**
  - ◆ Sending router no longer wants to be active router

## States



- **Initial**
- **Learn**
- **Listen**
- **Speak**
- **Standby**
- **Active**

# Virtual Router Redundancy Protocol (VRRP)

(C) Herbert Haas 2005/03/11

## Overview



- **RFC 3768 (Nokia 2004, VRRP v2)**
  - ◆ Uses Cisco patents!
- **Similar as HSRP but with the following basic differences:**
  - ◆ Master and Backup Routers
  - ◆ IP of master ("owner") is used instead of virtual IP
  - ◆ Hello interval greater or equal 1 second
    - Minimum failure detection time: 3 seconds!
  - ◆ Master has default priority of 255, backup has 100
- **VRRP-ID is used to cluster routers of the same redundancy group**

(C) Herbert Haas 2005/03/11

20

## Protocol



- **Virtual MAC address**
  - ♦ **00:00:5E:00:01:XX**
  - ♦ Where **XX** is the VRRP-ID
  - ♦ Note that **00:00:5E** is the OUI of the IETF
- **Multicast IP**
  - ♦ **224.0.0.18**

## Consequence of missing virtual IP



- **The IP address of the master is used similarly as a virtual node address**
  - ♦ If a backup router becomes active it will *respond to ARP requests for the owner-IP*
- **If a backup becomes active it will NOT respond to ping packets to the owner address**
  - ♦ *The backup only forwards packets with the virtual MAC-address and respond to ARP requests*

## Security



- **Current RFC (3768) removed authentication option!**

RFC 3768

Operational experience and further analysis determined that these did not provide any real measure of security. Due to the nature of the VRRP protocol, even if VRRP messages are cryptographically protected, it does not prevent hostile routers from behaving as if they are a VRRP master, creating multiple masters. Authentication of VRRP messages could have prevented a hostile router from causing all properly functioning routers from going into backup state. However, having multiple masters can cause as much disruption as no routers, which authentication cannot prevent. Also, even if a hostile router could not disrupt VRRP, it can disrupt ARP and create the same effect as having all routers go into backup.

## Configuration Example (HP Procurve)

```
(config)# router vrrp
(config)# vlan 54
(vlan54)# vrrp vrid 5
(vlan54)# ip address 1.2.3.4 /24
(vlan54)# owner !!! or 'backup'
(vlan54)# virtual-ip-address 1.2.3.4 /24 !!! must be same address !!!
(vlan54)# enable

# show vrrp vlan 54
# show vrrp config
```

# Common Address Redundancy Protocol (CARP)

(C) Herbert Haas 2005/03/11

## About



- **Created by OpenBSD team since patents prevent VRRP or HSRP implementations in open source projects**
  - ◆ See Linux ucarp package
- **Uses IP protocol number 112**
- **Tries to compensate VRRP and HSRP design flaws**
  - ◆ Strong cryptography (SHA-1 HMAC)
  - ◆ Low overhead
  - ◆ Supports also IPv6

(C) Herbert Haas 2005/03/11

26

# Gateway Load Balancing Protocol (GLBP)

(C) Herbert Haas 2005/03/11

## Overview



- Cisco proprietary
- Goal
  - ◆ Overcome with all limitations of HSRP and VRRP
  - ◆ Plus load balancing

(C) Herbert Haas 2005/03/11

28

## Proxy ARP

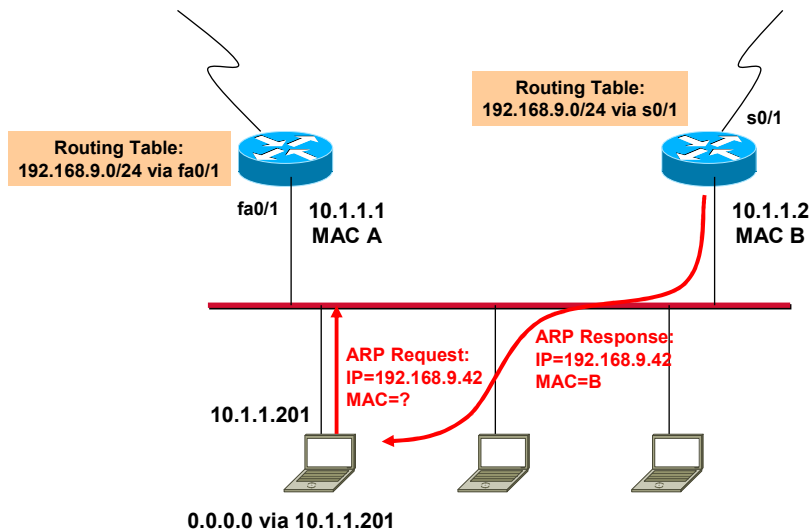
Old, simple, clever, sometimes  
useful, sometimes confusing

## Recall Proxy ARP



- On a host configure its own IP as default gateway
- Then the host performs an ARP request for any destination IP address
  - ♦ Even if the destination is in another network
  - ♦ The host does not need to know any official Default Gateway
- Cisco routers by default are Proxy-ARP enabled
  - ♦ That is, they will response to any ARP with their own MAC address – **if a route to the destination IP address is known**

## Example



(C) Herbert Haas 2005/03/11

31

## Proxy ARP + Mobile IP solution



- **Problem**
  - ♦ You have mobile devices (PCs or routers) with static IP address which are physically connected to one or the other backbone router
  - ♦ We want our backbone routers to detect such changes and inject host routes to the backbone routing tables
- **Solution**
  - ♦ Utilize Proxy ARP in combination with Mobile IP
  - ♦ Only the backbone routers must be configured

```
(config)# ip proxy-arp !!! by default enabled usually  
  
! Interface to mobile device  
(config-if)# mobile ip  
  
(config-router)# redistribute mobile
```

(C) Herbert Haas 2005/03/11

32

## Proxy ARP + Mobile IP solution



- The mobile device will try to determine the MAC of its home default router via ARP
- This is recognized by the foreign router which will provide its own MAC instead and redistribute a host route for that mobile client

## Problems with Proxy ARP



- No load balancing
- ARP cache might be invalid
  - ♦ One Default GW might fail, host continues to send packets to that dest-MAC
- Some newer IP hosts use the ICMP Router Discovery Protocol (IRDP) to find a new router when a route becomes unavailable
  - ♦ Can be abused by attackers to become MITM
  - ♦ Windows Vista gives IRDP higher precedence than DHCP (part of IPv6 stack)

# Service Level Agreement (SLA) Monitoring and Route Tracking

(C) Herbert Haas 2005/03/11

## About



- **Routers with multiple ISP connections should verify IP connectivity**
  - ◆ Link-up is not enough
- **Route tracking means that the reachabilities of particular IP destinations are verified periodically**
- **Typically done with ICMP echo request**
  - ◆ But also other possibilities, e. g. UDP/TCP-services
- **But also other network performance parameters should be checked**
  - ◆ Reachability is not the only requirement
  - ◆ Latency, drop probability, and jitter are also important!

(C) Herbert Haas 2005/03/11

36

## SLA Monitoring with Cisco Routers



- Cisco routers can simulate network data and IP services and collect network performance information in real time
  - ♦ Delay (one-way or RTT)
  - ♦ Jitter (interpacket delay variance)
  - ♦ Packet loss
  - ♦ Packet ordering
  - ♦ Voice quality scoring
  - ♦ Network resource availability (connectivity)
  - ♦ Application performance (download time)
- Either between Cisco IOS devices or from a Cisco IOS device to a remote IP device such as a network application server

## Configuration Steps



- Identify SLA source and responder
- Responder can be
  - ♦ Either a IOS device
    - Enabled via `ip sla responder` command
    - Configured via Cisco's SLA Control Protocol
  - ♦ Or any IP device (limited tests)
- Configure
  - ♦ Operation type (ICMP, UDP, HTTP, ...)
  - ♦ Options
  - ♦ Threshold conditions
  - ♦ Schedule execution (*now versus pending*)

## Note



- **Synchronize** both endpoints with the same time source via **NTP**
  - ◆ Needed for delay measurements
  - ◆ Not needed for one-way jitter measurements
- **SLA results are accessible via SNMP**
- **SLA responder functionality is not required for certain operation types (e. g. ICMP-based)**
- **Optionally specify the IP address and port on the responder for which it should listen to (=local address!)**

```
(config)# ip sla responder

(config)# ip sla responder tcp-connect ipaddress <address>
                                     port <port-number>

!!! Required for UDP jitter operations:
(config)# ip sla responder udp-echo ipaddress <address>
                                     port <port-number>
```

## SLA Operation Types



- DLSW+ (Data Link Switching Plus)
- DNS
- DHCP
- FTP
- HTTP
- ICMP Echo
- ICMP jitter
- ICMP Path Echo
- ICMP Path Jitter
- RTP-based VoIP
- TCP Connect
- UDP Echo
- UDP jitter
- UDP jitter for VoIP
- VoIP gatekeeper registration delay
- VoIP post-dial delay

## SLA Thresholds



- An SLAs threshold violation can trigger
  - ♦ An SNMP trap
  - ♦ Another SLAs operation
    - For example, the frequency could be increased or an ICMP path echo or ICMP path jitter operation could be initiated for troubleshooting
- **threshold <msec>** command
  - ♦ For **UDP jitter operation** specifies upper *jitter* threshold
  - ♦ For **all other operations** it specified the upper threshold value for the *RTT*
- **ip sla reaction-configuration** command
  - ♦ Allows to specify **other thresholds types**
  - ♦ To trigger SNMP traps or start another SLA operation

## UDP Jitter Operation



- **Not only measures jitter, but:**
  - ♦ Per-direction jitter (source to destination and destination to source)
  - ♦ Per-direction packet-loss
  - ♦ Per-direction delay (one-way delay)
  - ♦ Round-trip delay (average round-trip time)
- **Requires SLA Responder**

## UDP Jitter Operation (cont.)



- **Most precise jitter measurement** compared to other SLA methods
- By default a burst of 10 packets are sent every minute
  - ◆ Inter-packet delay is 20 ms
  - ◆ Packet size is 60 bytes total (IP+UDP+Payload = 20+8+32 bytes)

## UDP Jitter Operation - Configuration



```
(config)# ip sla <sla-number>
(config-ip-sla)# udp-jitter <dest-ip> <dest-port>
                        [source-ip <src-ip>]
                        [source-port <src-port>]
                        [control enable|disable]
                        [num-packets <num-packets>]
                        [interval <msec>]
(config-ip-sla-jitter)# request-data-size <bytes>
(config-ip-sla-jitter)# tos 128 !!! whole byte in decimal
(config-ip-sla-jitter)# frequency <seconds>

(config)# ip sla schedule <sla-number>
                        [life forever|<seconds>]
                        [start-time <...>]
                        [ageout <sec>]
                        [recurring]
```

## UDP Jitter Operations – Example



```
ip sla 1
  udp-jitter 192.168.0.2 42525 num-packets 20
  frequency 120
ip sla schedule 1 life 300 start-time after 00:05:00
```

## ICMP Jitter Operation



- **Measures**
  - ♦ Jitter (source-to-destination and destination-to-source)
  - ♦ Latency (source-to-destination and destination-to-source)
  - ♦ Round-trip time latency
  - ♦ Packet loss
  - ♦ Successive packet loss
  - ♦ Out-of-sequence packets (source-to-destination, destination-to-source, and round-trip)
  - ♦ Late packets
- **Less accurate than UDP method**
- **Destination can be any IP host – no SLA responder feature required**

## ICMP Jitter Operation – Example



```
ip sla 42
 icmp-jitter 192.168.1.2 [interval <ms> num-packets <num>]
 timeout 1000
 threshold 500 !!! RTR
 frequency 300 !!! Repeat every 5 minutes
 history hours-of-statistics-kept 1 !!! keep data 1 hour
```

```
ip sla 1
 icmp-jitter 192.168.3.2 interval 40 num-packets 100
                                     source-ip 10.42.42.1
 frequency 50
 !
 !!! config mode now:
 ip sla reaction-configuration 1 react jitterAvg threshold-value 5 2
                                     action-type trap
 !
 ip sla schedule 1 start-time now life forever
```

## ICMP Path Jitter Operation



- **Detects bottlenecks along the path**
  - ◆ First discovers each hop using traceroute
  - ◆ Then measures statistics to each hop
- **Cannot compensate for destination processing delay**
  - ◆ As UDP Jitter operation does using two timestamps
- **SLA responder feature NOT required**

## ICMP Path Jitter - Example



```
(config)# ip sla 10
(config-ip-sla)# path-jitter 192.168.42.1 source-ip 10.1.1.1
num-packets 20
```

```
# show ip sla statistics

Current Operational State
Entry Number: 10
Modification Time: 23:08:31.122 UTC Thu Jul 18 2007
Diagnostics Text:
Last Time this Entry was Reset: Never
Number of Octets in use by this Entry: 1885
Number of Operations Attempted: 1
Current Seconds Left in Life: 3466
Operational State of Entry: active
Latest Completion Time Average (milliseconds): 4
Latest Operation Start Time: 03:20:21.314 UTC Thu Jul 18 2007
Path Jitter Statistics:
Legend - TR = Total Receives; RTT = Round Trip Time (Avg); PL = Packet Loss;
        DS = Discarded Samples; OS = Out Of Sequence Echo Replies
HopAddress      TR    RTT    PL    DS    OS    Jitter(RFC 1889)
10.1.1.2         10     1     0     0     0     0
172.16.20.1      10     1     0     0     0     0
172.17.112.122   10     1     0     0     0     0
171.31.4.16      10     1     0     0     0     0
171.31.5.6       10     1     0     0     0     0
```

## RTP-based VoIP Operation



- Establishes a test call and measures DSP quality
- Measures
  - ◆ Jitter
  - ◆ Frame Loss
  - ◆ MOS for Conversational Quality (MOS-CQ)
  - ◆ MOS for Listening Quality (MOS-LQ)

## Requirements



- **Both sides must use**
  - ♦ IOS Release 12.4(4)T or later
  - ♦ IOS IP Voice or higher grade feature package
- **The source router must have a network module with a c5510 or c549 DSP**
  - ♦ The destination router need NOT have a network module with a DSP
  - ♦ Statistics only gathered from DSP of source router
- **SLAs Responder must be enabled at the destination**

## Configuration Example



```
ip sla 1
  voip rtp 10.2.2.2
    source-ip 10.1.1.1
    source-voice-port 1/0:1
    codec g711alaw
    duration 30          !!! duration of the test call (sec)
!
ip sla reaction-configuration 1 react FrameLossDS action-type
  traponly threshold-type consecutive 3
!
ip sla schedule 1 start-time now life forever
```

## Verification



```
!!! Check which SLA operations are available on this IOS:
# show ip sla application

!!! Check configuration
# show ip sla configuration [sla-number]

# show ip sla statistics aggregated
```

## Verification (cont.)



```
# show ip sla statistics aggregated

Round Trip Time (RTT) for      Index 50000
Start Time Index: 20:48:12.051 UTC Thu May 22 2009
Type of operation: icmpJitter
RTT Values:
    Number Of RTT: 281          RTT Min/Avg/Max: 5/17/43
Latency one-way time:
    Number of Latency one-way Samples: 0
    Source to Destination Latency one way Min/Avg/Max: 0/0/0
    Destination to Source Latency one way Min/Avg/Max: 0/0/0
Jitter Time:
    Number of Jitter Samples: 225
    Source to Destination Jitter Min/Avg/Max: 1/7/25
    Destination to Source Jitter Min/Avg/Max: 1/1/1
```