



Hot Standby Routing Protocol Virtual Router Redundancy Protocol

Pietro Nicoletti
piero[at]studioreti.it



Copyright note

This set of transparencies, hereinafter referred to as slides, is protected by copyright laws and provisions of International Treaties. The title and copyright regarding the slides (including, but not limited to, each and every image, photography, animation, video, audio, music and text) are property of the authors specified on page 1.

The slides may be reproduced and used freely by research institutes, schools and Universities for non-profit, institutional purposes. In such cases, no authorization is requested.

Any total or partial use or reproduction (including, but not limited to, reproduction on magnetic media, computer networks, and printed reproduction) is forbidden, unless explicitly authorized by the authors by means of written license.

Information included in these slides is deemed as accurate at the date of publication. Such information is supplied for merely educational purposes and may not be used in designing systems, products, networks, etc. In any case, these slides are subject to changes without any previous notice. The authors do not assume any responsibility for the contents of these slides (including, but not limited to, accuracy, completeness, enforceability, updated-ness of information hereinafter provided).

In any case, accordance with information hereinafter included must not be declared.

In any case, this copyright notice must never be removed and must be reported even in partial uses.



Default gateway redundancy

■ Scope:

- mission critical network need redundant default gateway
- modern solution for hosts (substitute route demon old solution)

■ Solutions:

- HSRP (Hot Standby Routing Protocol): Cisco proprietary protocol defined on RFC 2281 (informational type)
- VRRP (Virtual Router Redundancy Protocol): standard protocol defined on RFC 3768 (2338 old)
- IPSTB (IP Standby protocol): Digital Equipment Corporation proprietary protocol




HSRP: general aspects

- HSRP packets:
 - UDP encapsulation using 1985 port
 - Transmitted to multicast address 224.0.0.2
- TTL = 1 (only one possible hop)
- Possibility to configure more HSRP groups
 - Each standby group emulates a single virtual router
 - For each standby group, a single well-known MAC address is allocated to the group, as well as an IP address
 - Load balancing can be achieved by distributing hosts among different standby groups



HSRP: overlapping groups and load balancing

- Overlapping groups on a router interface
 - For any of overlapped group the router interface can become Active or Standby according to the priority that has been configured
 - Load balancing over different group



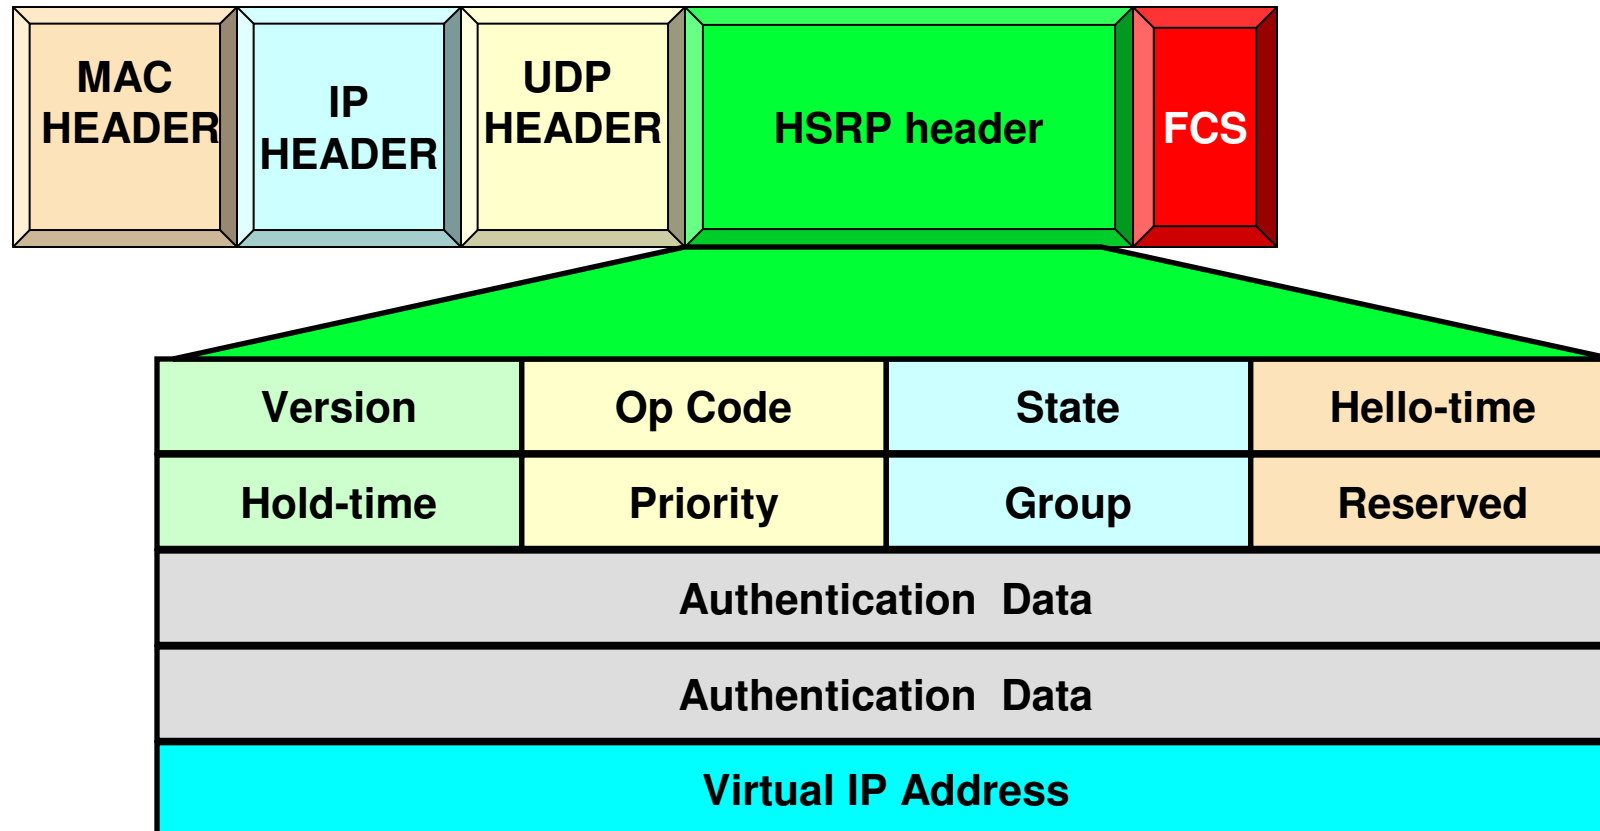


HSRP functions

- A single router elected from the group is responsible for forwarding the packets that hosts send to the virtual router
- The *Active* router interface has the followings assigned addresses:
 - Primary IP address (inserted in the source IP header field)
 - Virtual IP address (used by hosts as default gateway)
 - Well-Known MAC address allocated to the HRSP Group
- The *Standby* router interface has the followings assigned addresses:
 - Primary IP address (inserted in the source IP header field)
 - Physical MAC address assigned by the producer (unique in the world)
- Both interfaces: Active and Standby send hello messages containing different information



HSRP packet format





HSRP header: OP Code field

- Describes the type of message contained in this packet; possible values are:
 - 0 = Hello
 - indicate that a router is running and is capable of becoming the active or standby router
 - 1 = Coup
 - coup messages are sent when a router wishes to become the active router
 - 2 = Resign
 - resign messages are sent when a router no longer wishes to be the active router



HSRP header: State field (1st part)

- Describes the current state of the router sending the message; possible values are:
 - 0 = Initial
 - This is the starting state and indicates that HSRP is not running
 - 1 = Learn
 - The router has not determined the virtual IP address and is still waiting to hear from the active router
 - 2 = Listen
 - The router knows the virtual IP address, but is neither the active router nor the standby router



HSRP header: State field (2nd part)

■ State field:

■ 4 = Speak

- The router sends periodic Hello messages and is actively participating in the election of the active and/or standby router

■ 8 = Standby

- The router is a candidate to become the next active router and sends periodic Hello messages
- Must be at most one router in the group in Standby state

■ 16 = Active

- The router is currently forwarding packets that are sent to the group's virtual MAC address
- Must be at most one router in Active state in the group



HSRP header: Hello and Hold time fields

■ Hello-Time:

- Period between the Hello messages that the router sends
- If the Hello-time is not learned from a Hello message from the active router and it is not manually configured, a default value of 3 seconds is used

■ Hold-Time:

- Amount of time that the current Hello message should be considered valid.
- When this timer expires the Standby router proposes him as Active router
- If the Hold-time is not learned and it is not manually configured, a default value of 10 seconds is used



HSRP header: Priority and Group fields

■ Priority:

- Used to elect the active and standby routers
- The router with the numerically higher priority wins and become Active.
- In the case of routers with equal priority the router with the higher IP address wins
- Default Priority value is 100

■ Group:

- For Token Ring, values between 0 and 2 inclusive are valid
- For other media values between 0 and 255 inclusive are valid



HSRP header: Authentication e Virtual IP Address

■ Authentication Data:

- This field contains a clear-text 8 character reused password
- If no authentication data is configured the default text is "cisco"

■ Virtual IP address:

- Used by hosts as default gateway IP address
- The virtual IP address used by this group
- If the virtual IP address is not configured on a router, then it may be learned from the Hello message from the active router.



Preemption capability

- If a router has higher priority than the active router and preemption is configured, it MAY take over as the active router using a Coup message
 - Without preemption function configured the higher priority can't force the router to become Active



HSRP basic configuration

- How to configure HSRP parameters at interface level:

```
interface Ethernet0
```


```
ip address 10.1.1.1 255.255.255.0
```

```
standby 24 preempt
```

```
standby 24 ip 10.1.1.5
```

```
standby 24 priority 105
```

- 24 is the HSRP Group number in this example





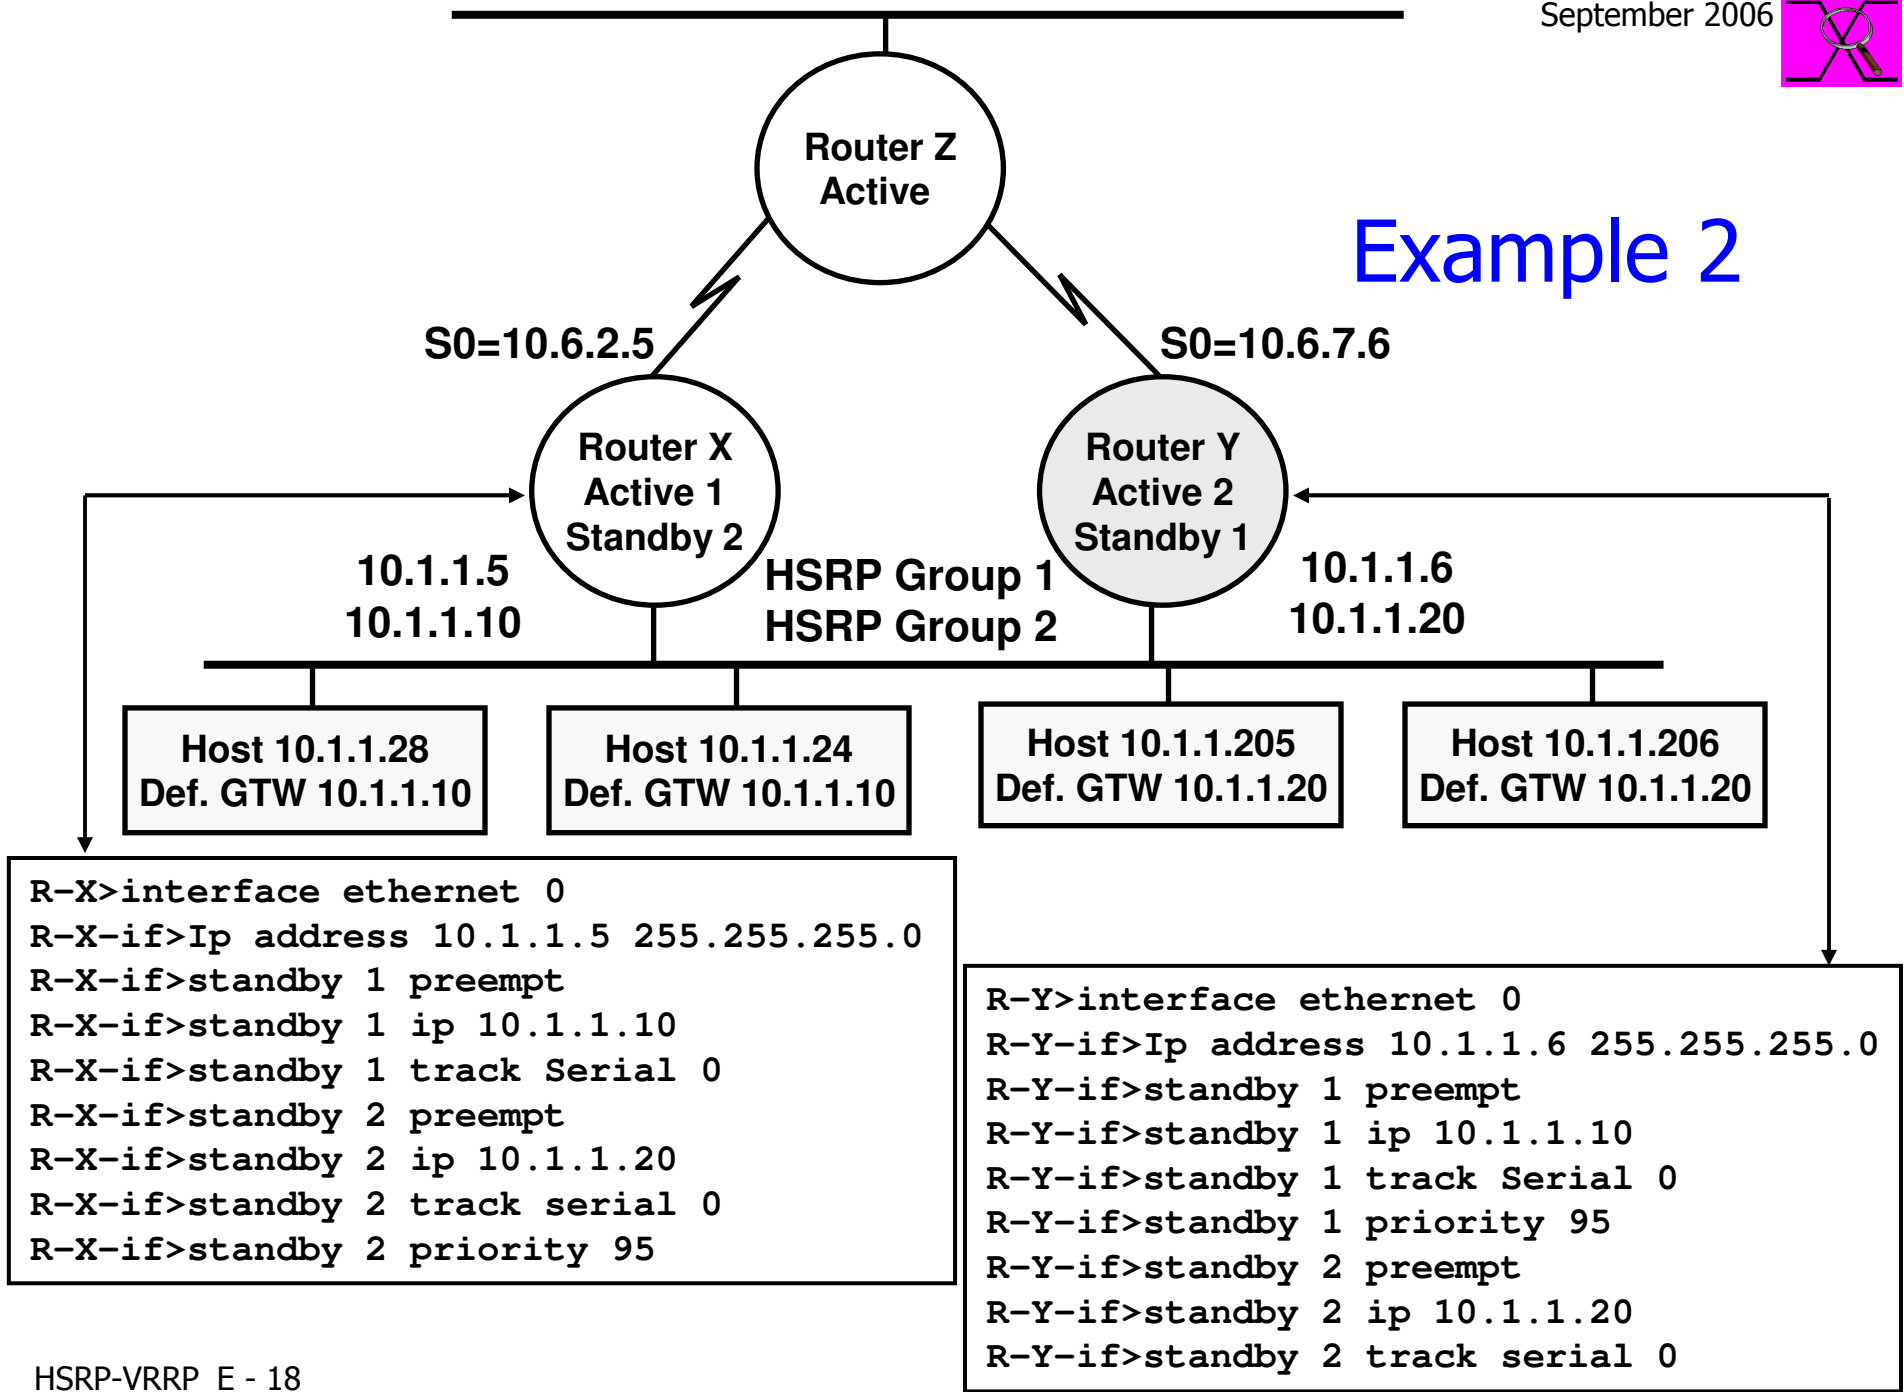
Cisco router Track function

- This function dynamically decrease the HSRP Priority when a tracked interface goes down.
 - By default HSRP algorithm decrease the Priority by 10 when a tracked interface goes down.
 - Be care full! Not any fault in a connection cause the interface down.

```
standby 1 preempt  
standby 1 ip 10.1.1.10  
standby 1 track Serial 0
```



Example 2





HSRP: Well-Known MAC address

- Act as a Virtual MAC Address
- Take in charge by the Active router and used to answer at the received ARP request packets
- Token Ring well known virtual MAC address corresponding to HSRP Group 0, 1, 2
 - C0-00-00-01-00-00
 - C0-00-00-02-00-00
 - C0-00-00-04-00-00
- Well known virtual MAC address for other LAN type (example 802.3, 802.11 etc.)
 - 00-00-0C-07-AC-XX
 - XX It represents the HSRP Group



VRRP: general aspects (1st part)

- VRRP packet:
 - IP encapsulation
 - Transmitted to multicast address **224.0.0.18**
 - TTL = **255**
 - A VRRP router receiving a packet with the TTL not equal to 255 must discard the packet (only one possible hop)
- A VRRP Router may backup one or more virtual routers.
- Any of the virtual router's IP addresses on a LAN can then be used as the Default Gateway by end-hosts.
- Master VRRP router may control one or more IP Addresses.



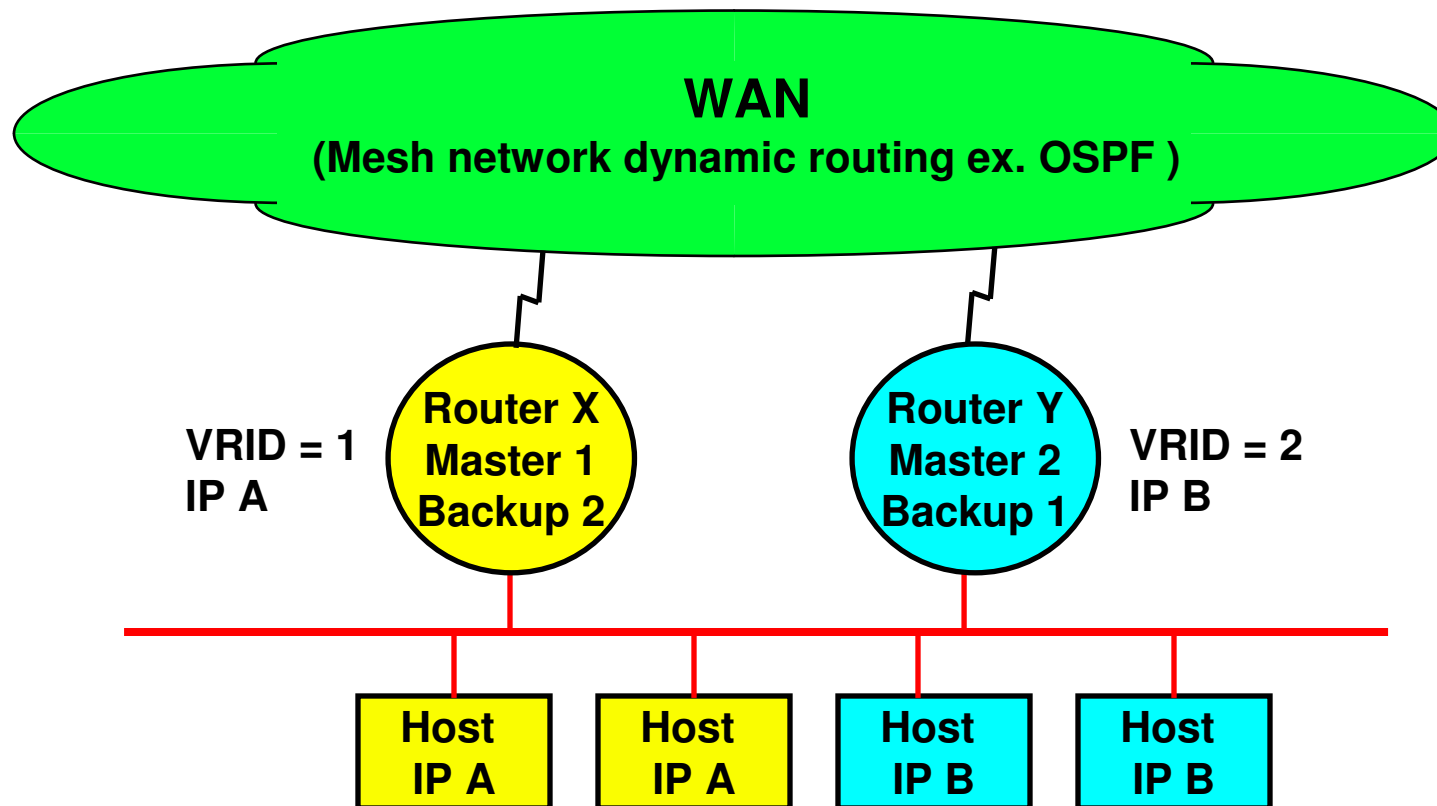
VRRP: general aspects (2nd part)

- Support of multiple logical IP subnets on a single LAN segment
- Virtual Router: concepts similar as in HSRP:
 - VRID (Virtual Router Identifier) identify a Virtual Router
 - as well HSRP group emulate a Virtual router
- For any VRID a single Master Router is elected the remaining routers are selected as Backup Router



Election of multiple virtual routers on a network for load balancing

- Load balancing can be achieved by distributing hosts among different VRIDs



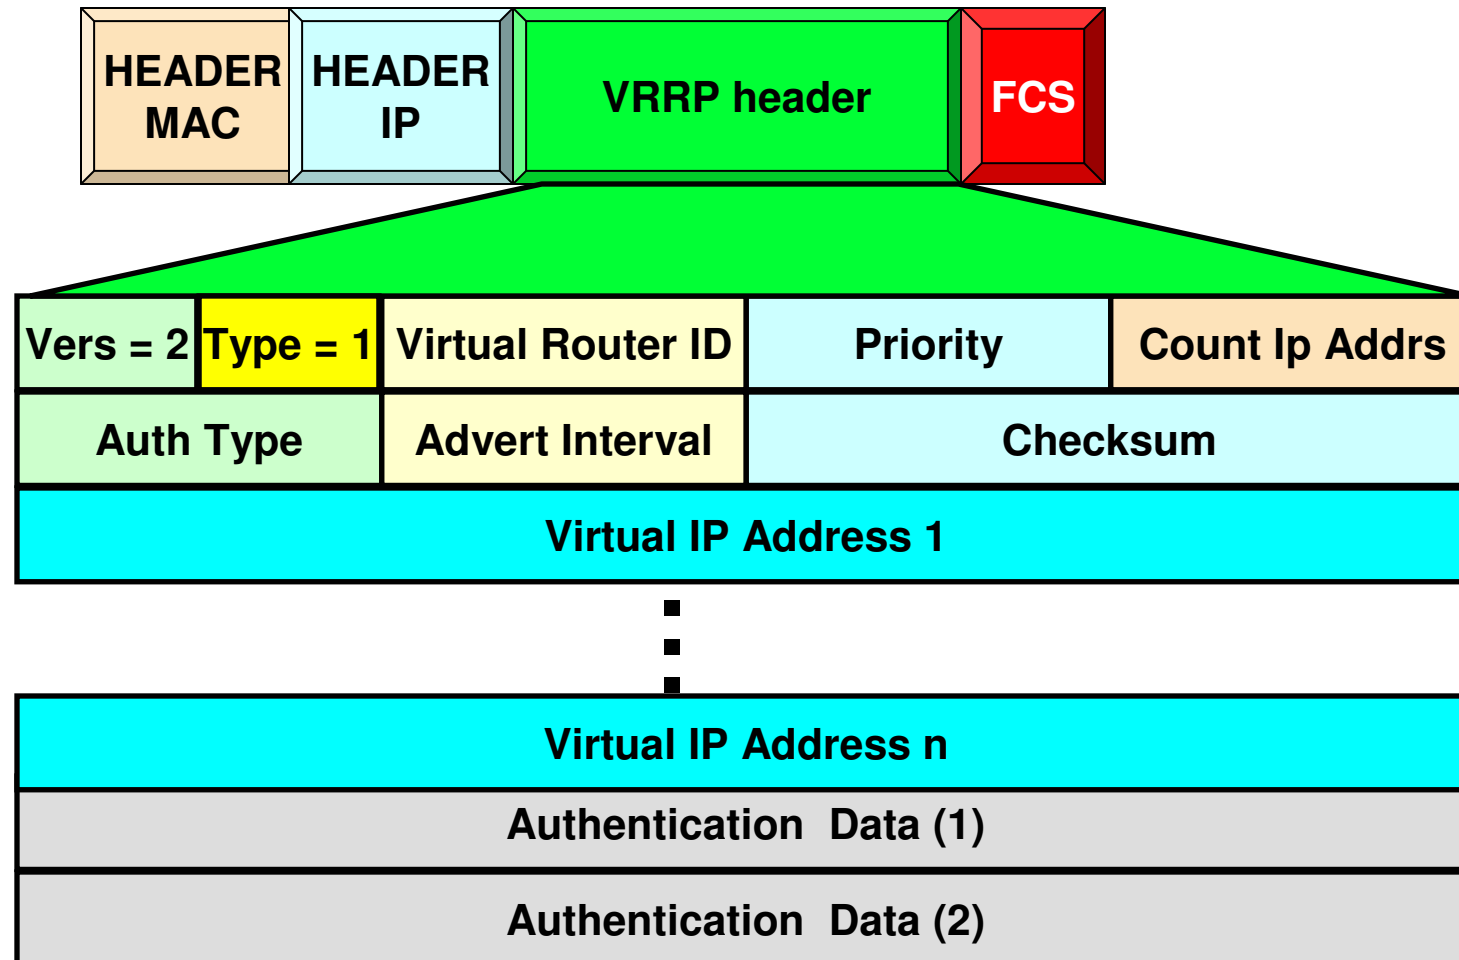


VRRP: functions

- The VRRP router controlling the IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses
- The **Master** router has the followings assigned addresses:
 - Primary IP address (inserted in the source IP header field)
 - One or more Virtual IP addresses (used by hosts as default gateway)
 - Well-Known MAC address allocated to VRID
- The **Backup** router has the followings assigned addresses:
 - Primary IP address (inserted in the source IP header field)
 - Physical MAC address assigned by the producer (unique in the world)
- Only the Master router send Advertisement packets



VRRP: packet format





VRRP packet and IP Header

- The TTL MUST be set to 255. A VRRP router receiving a packet with the TTL not equal to 255 MUST discard the packet
- The IP protocol number assigned by the IANA for VRRP is 112 (decimal)



VRRP Field Descriptions 1st part

■ Type

- The type field specifies the type of this VRRP packet. The only packet type is:
 - 1 Advertisement

■ VRID

- The Virtual Router Identifier (VRID) field identifies the virtual router this packet is reporting status for



VRRP Field Descriptions 2nd part

■ Priority

- The priority value for the VRRP router that owns the IP address(es) associated with the virtual router must be 255
- VRRP routers backing up a virtual router **MUST** use priority values between 1-254
- The default priority value for VRRP routers backing up a virtual router is 100
- The priority value zero (0) has special meaning indicating that the current Master has stopped participating in VRRP

■ Preempt_Mode

- Controls whether a higher priority Backup router preempts a lower priority Master



VRRP Field Descriptions 3rd part

- Count IP Adrs
 - number of IP addresses contained in this VRRP Advertisement
- Authentication Type
 - 0 = No Authentication
 - 1 = Simple Text Password
 - 2 = IP Authentication Header
 - HMAC-MD5-96 within ESP and AH" [HMAC]



VRRP Timer

- Advertisement Interval
 - The Advertisement interval indicates the time interval (in seconds) between Advertisements
 - default value = 1 s
- Skew_Time
 - $(256 - \text{Priority}) / 256$
- Master_Down_Interval
 - $(3 * \text{Advertisement_Interval}) + \text{Skew_time}$
 - Time interval for Backup to declare Master down (seconds)



VRRP: Token Ring well known virtual MAC address

VRID	Token Ring Functional Address
1	03-00-02-00-00-00
2	03-00-04-00-00-00
3	03-00-08-00-00-00
4	03-00-10-00-00-00
5	03-00-20-00-00-00
6	03-00-40-00-00-00
7	03-00-80-00-00-00
8	03-00-00-01-00-00
9	03-00-00-02-00-00
10	03-00-00-04-00-00
11	03-00-00-08-00-00



VRRP: Virtual MAC Address

- Well known virtual MAC address for other LAN type (example 802.3, 802.11 etc.)
 - 00-00-5E-00-01-XX
 - XX It represents the VRID



VRRP configuration example on Alcatel 8800 Layer 3 switch

- IP Address and VRRP configuration for VLAN 1
 - The Advertisement Interval value as been increased to 20 second to work fine with STP convergence

- Configuration:

```
! VLAN :
```

```
vlan 1 enable name "Default"
```

```
vlan 1 router ip 172.14.0.253 255.255.255.0 e2
```

```
!
```

```
.....
```

```
! VRRP :
```

```
VRRP 1 1 PRIORITY 110 PREEMPT INTERVAL 20 AUTHENTICATE test-1
```

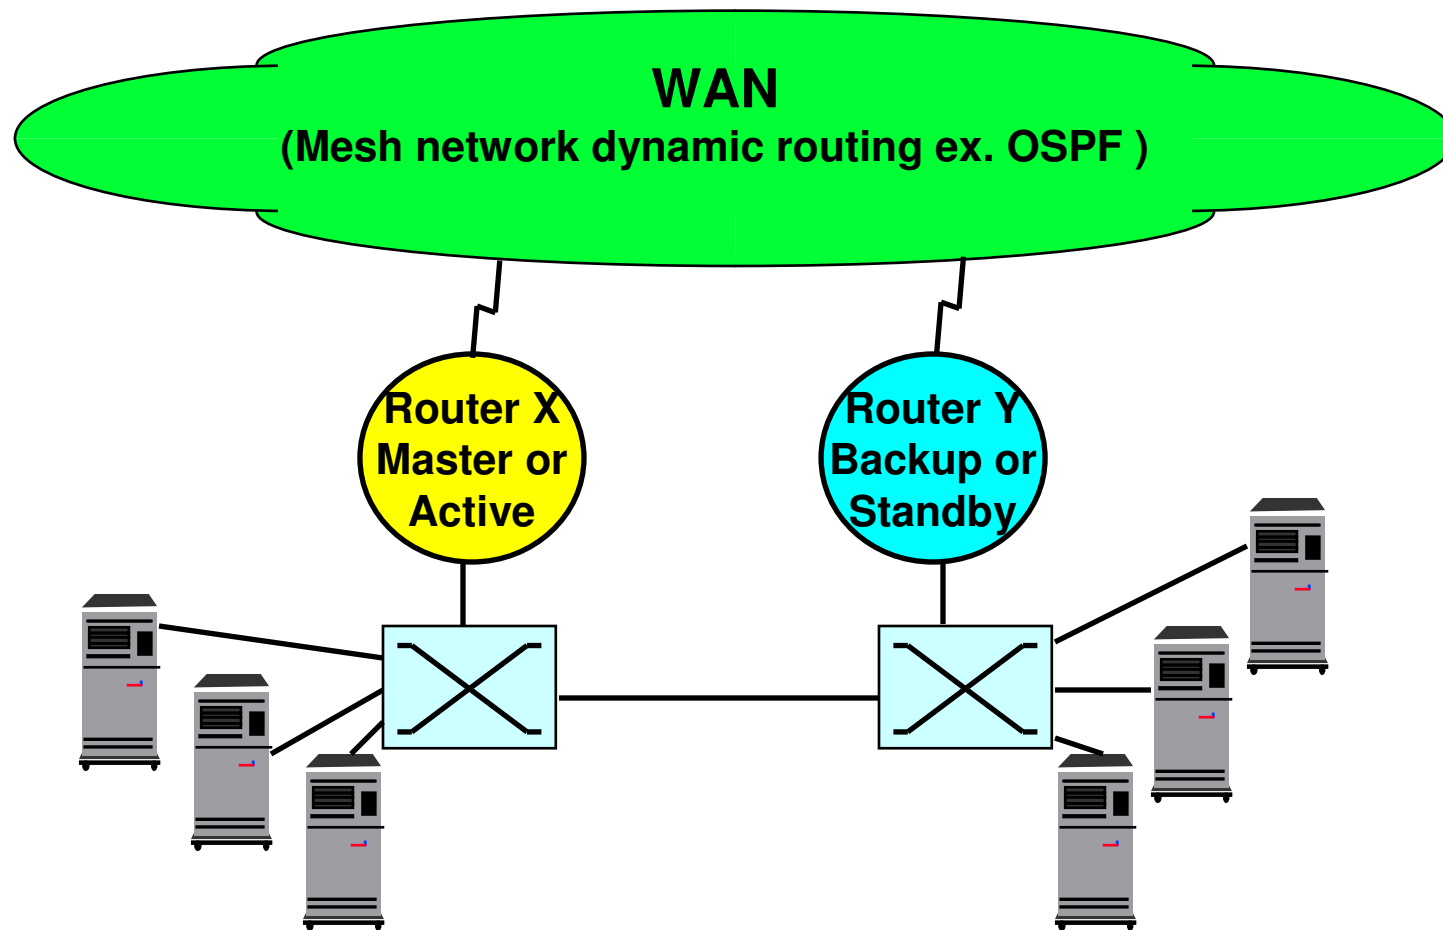
```
VRRP 1 1 IP 172.14.0.254
```

```
VRRP 1 1 ENABLE
```




HSRP/VRRP over Switched LAN 1st part


- Wrong architecture





HSRP/VRRP over Switched LAN 2nd part


- 1st case of fault: OK





HSRP/VRRP over Switched LAN 3rd part


- 2nd case of fault: OK





HSRP/VRRP over Switched LAN 4th part

- 3rd case of fault: same logical IP subnet over different physical networks!!!





HSRP/VRRP over Switched LAN: the solution

- Correct architecture

