kenet
Kenya Education Network

HIGH AVAILABILITY

Presented by

MICHELLE OPIYO

Transforming education using ICT

HIGH AVAILABILITY

- How can we achieve high availability?
- By protecting your network against a single device failure affecting all of your network
- By introducing hardware resiliency and backup paths
- Different techniques depending on the layer
- There is a relationship between reliability, complexity and cost
- The trick is to balance all variables and come up ahead

HIGH AVAILABILITY

- You need to evaluate your needs
- **Minimal need**
  - Network just needs to be up for a portion of the day
  - Downtime is easily scheduled after working hours
  - Business is not impacted if the network is down
  - Users' productivity is not impacted by a network failure

HIGH AVAILABILITY

- **Medium need**
  - Network needs to be available for most of the day
  - Only centralized servers need to be up 24 hours/day
  - Downtime needs to be scheduled on weekends
  - If critical parts of the network fail, the business operation is impacted
  - A network failure affects user productivity

HIGH AVAILABILITY

- **High need**
- Network needs to be up 24x7
- Downtime needs to be scheduled well in advance and completed within schedule
- A network failure causes major loss of business
- User productivity drastically impacted by a network failure

HIGH AVAILABILITY

- **METHODS**
- Component Redundancy
- Duplicate or backup parts
- Power supplies, fans, processors, etc

HIGH AVAILABILITY

- Server Redundancy
- Protect your data with backups
- Use of hot standby servers
- Use of load balancers

HIGH AVAILABILITY

- Network Link & Data Path Redundancy
- Provide physical redundant connections between devices
- Allow for hot backup paths (STP) and parallelism (routing)

HIGH AVAILABILITY

- So I built all this redundancy and high availability in my network, how can my end users take advantage of it?
- You are already providing more that one router for each subnet
- You want to provide your users with a way to move their traffic from one default gateway to another

First-hop redundancy protocals

- HSRP
- VRRP
- GLBP

HSRP

- Cisco proprietary protocol
- Allows several routers/layer 3 switches to appear as a single gateway IP address

HSRP

- One router is elected as primary and the other as the standby.
- All the other routers will remain in the listening state.
- HSRP sends it hello messages to the multicast destination 224.0.0.2 using UDP port 1985

HSRP

- A HSRP group can be assigned a group number from 0 to 255.
- Most switches support up to 16 unique group numbers.
- The group numbers however are unique to each VLAN/ Interface. So group 1 on VLAN 11 is unique to group 1 on VLAN 111

HSRP

- switch(config-if)# **standby** group **priority** priority
- switch(config-if)# **standby** 2 **priority** 200

HSRP

- HSRP states
  - Disabled
  - Init
  - Listen
  - Speak
  - Standby
  - Active

HSRP

- Only the standby router (2nd highest priority) receives hellos from the the active router
- Default hello timer 3 secs, holdtime 10 secs
- **switch(config-if)# standby 2 timers msec 10 msec 30**
- The timers can be adjusted. Decreasing the timers means router failure is detected faster. However, more hellos increases amount of traffic on the interface.

HSRP

- When the active router fails and the standby becomes active, the original active router cannot immediately become active once it is restored even with a higher priority
- You can configure the router with the highest priority to preempt.
- **switch(config-if)# standby group preempt [delay [min secs] [reload sec]]**

HSRP

- switch(config-if)# **standby** group **ip** ip-address [**secondary**]
- BB11(config)# int vlan 64
- BB11(config-if)# ip add 10.1.64.4 255.255.255.0
- BB11(config-if)# standby 1 priority 200
- BB11(config-if)# standby 1 preempt
- BB11(config-if)# standby 1 ip 10.1.64.10

HSRP Configuration

- switch(config-if)# **standby** group **ip** ip-address [**secondary**]

- BB12(config)# int vlan 64
- BB12(config-if)# ip add 10.1.64.5 255.255.255.0
- BB12(config-if)# standby 1 priority 100
- BB12(config-if)# standby 1 preempt
- BB12(config-if)# standby 1 ip 10.1.64.10

HSRP Configuration

- The Virtual IP address is associated with a special mac address in the form 0000.0c0.acXX
- XX represents the HSRP group number in Hexadecimal value.
- eg. for HSRP group 1 the mac address will be 0000.0c07.ac01

- In a network in which HSRP has been deployed on 2 distribution switches, one becomes active and the other standby.
- In order to load balance across the 2 links, configure 2 HSRP groups with one of the layer 3 switches as active on each group and standby by for the other.
- Configure the end user's default gateway as 1 of the 2 virtual HSRP group addresses.

HSRP LOAD BALANCING

- B11(config)# int vlan 64
- BB11(config-if)# ip add 10.1.64.4 255.255.255.0
- BB11(config-if)# standby 1 priority 200
- BB11(config-if)# standby 1 preempt
- BB11(config-if)# standby 1 ip 10.1.64.10
- BB11(config-if)# standby 2 priority 100
- BB11(config-if)# standby 2 preempt
- BB11(config-if)# standby 2 ip 10.1.64.11

HSRP LOAD BALANCING

- BB12(config)# int vlan 64
- BB12(config-if)# ip add 10.1.64.5 255.255.255.0
- BB12(config-if)# standby 1 priority 100
- BB12(config-if)# standby 1 preempt
- BB12(config-if)# standby 1 ip 10.1.64.10
- BB12(config-if)# standby 2 priority 200
- BB12(config-if)# standby 2 preempt

HSRP VERIFICATION

- Switch# show standby [brief] vlan vlan-id | type mod/num]

- BB11# show standby vlan 64 brief
- BB11# show standby vlan 64

VRRP

- Virtual Router Redundancy Protocol (VRRP) is the standards based alternative to HSRP defined in RFC 2338.
- Similar to HSRP with slightly different terminology and slight functional differences

VRRP

- Group numbers range from 0 – 255
- Priorities range from 1 – 254 (100 is the default)
- Virtual mac address 0000.5e00.01XX (where XX is the 2 digit hex VRRP group number)
- By default VRRP will preempt the current master if the priority is greater.
- VRRP uses mutlicast address 224.0.0.18

VRRP LOAD BALANCING

- B11(config)# int vlan 64
- BB11(config-if)# ip add 10.1.64.4 255.255.255.0
- BB11(config-if)# vrrp 1 priority 200
- BB11(config-if)# no vrrp 1 preempt
- BB11(config-if)# vrrp 1 ip 10.1.64.10
- BB11(config-if)# vrrp 2 priority 100
- BB11(config-if)# no vrrp 2 preempt
- BB11(config-if)# vrrp 2 ip 10.1.64.11

VRRP LOAD BALANCING

- BB12(config)# int vlan 64
- BB12(config-if)# ip add 10.1.64.5 255.255.255.0
- BB12(config-if)# vrrp 1 priority 100
- BB12(config-if)# no vrrp 1 preempt
- BB12(config-if)# vrrp 1 ip 10.1.64.10
- BB12(config-if)# vrrp 2 priority 200
- BB12(config-if)# no vrrp 2 preempt
- BB12(config-if)# vrrp 2 ip 10.1.64.11

VRRP VERIFICATION

- Switch# show VRRP [brief]

GBLP

- A Cisco proprietary protocol designed to overcome the limitations of existing redundant router protocols
- To provide a virtual gateway, multiple switches are assigned a common GLBP group. Instead of having just one router performing the forwarding for the virtual router address, all routers in the group can participate and offer load balancing

GLBP

- All the end users can have the same default gateway
- Load balancing is achieved through the use of a virtual mac address in ARP replies returned to the clients.
- The client request an ARP and the AVG (active virtual gateway) responds with a virtual mac for on of the routers in the group.
- Clients end up using same gateway IP but different mac addresses.

GLBP

- **Active Virtual Gateway (AVG)**
- One of the routers is elected as the AVG.
- The router with the highest priority is selected else the router with the highest IP in the group will be selected.
- The mac address it returns is based on the load balancing algorithm used.

GLBP

- The AVG assigns the mac addresses to the participating routers in the group (up to 4 routers) .
- The forwarding routers in the group are referred to as an AVF (Active Virtual Forwarder)
- The virtual mac address is in the form 0007.b4xx.xxyy
- Where xx.xx represents six 0 bits followed by the GLBP group number. The 8-bit yy value is the virtual forwarder number

- Periodic hellos are used to detect AVF failures. If hellos are not received from a peer before its Hold-Time timer expires, the AVG assumes the AVF has failed and assigns the AVF role to another router.
- At this point, clients using the old mac need to refresh their arp cache to obtain the new virtual mac address.

GLBP LOAD BALANCING

- The AVG establishes load balancing by handing out virtual mac addresses in a deterministic manner.
- **Round Robin –** Each new ARP request for virtual router address receives the next available mac
- **Weighted –** The GLBP interface weighting value determines proportion of traffic that should be sent to that AVF
- **Host dependent –** The host always receives the same virtual mac address is arp reply.

GLBP CONFIGURATION

- BB11(config)# int vlan 64
- BB11(config-if)# ip add 10.1.64.4 255.255.255.0
- BB11(config-if)# glbp 1 priority 200
- BB11(config-if)# glbp 1 preempt
- BB11(config-if)# glbp 1 ip 10.1.64.10

GLBP CONFIGURATION

- BB11(config)# int vlan 64
- BB11(config-if)# ip add 10.1.64.4 255.255.255.0
- BB11(config-if)# glbp 1 priority 200
- BB11(config-if)# glbp 1 preempt
- BB11(config-if)# glbp 1 ip 10.1.64.10

GLBP CONFIGURATION

- BB12(config)# int vlan 64
- BB12(config-if)# ip add 10.1.64.5 255.255.255.0
- BB12(config-if)# glbp 1 priority 100
- BB12(config-if)# glbp 1 preempt
- BB12(config-if)# glbp 1 ip 10.1.64.10

GLBP VERIFICATION

- sh glbp [brief]

CONCLUSION

All this redundancy and high availability is not going to do you any good if:

- You don’t test it
- Make sure that it actually works the way you expect
- You don’t monitor it
- If the redundant devices or links are down, it won’t work!

kenet
Kenya Education Network

Transforming education using ICT

Thank You

www.kenet.or.ke
Jomo Kenyatta Memorial
Library, University of Nairobi
P. O Box 30244-00100, Nairobi.
0732 150 500 / 0703 044 500