AlliedWare™ OS

How To | Configure VRRP (Virtual Router Redundancy Protocol)

Introduction

VRRP is a popular protocol for providing device redundancy, for connecting redundant WAN gateway routers or server access switches. It allows a backup router or switch to automatically take over if the primary (master) router or switch fails.

This How To Note describes one possible basic VRRP configuration.

VRRP works by grouping the redundant routers together into a single *virtual router*. That virtual router entity has an IP address of its own. Instead of sending traffic to an individual router, PCs etc send traffic to the virtual router address (for example, by using the virtual router address as their gateway address). The master router processes traffic that is addressed to the virtual router address and forwards it appropriately.

The master router also sends out regular advertisements to the backup router. If the master router goes down, the backup router stops receiving these advertisements. In that case, the backup router takes over as the master router and starts processing traffic. When the original master router comes back up, it takes over as the master router again.

For more information about VRRP, see the Software Reference for your router or switch.

Which products does this note apply to?

This configuration applies to the following Allied Telesis switches, running all supported AlliedWare software versions:

- AT-8600, AT-8700XL, AT-8800, Rapier, Rapier w, and Rapier i series
- AT-8948, AT-9900, AT-9900s, and x900 series
- SwitchBlade 4000 and AT-9800 series

VRRP is also available on AR400 and AR700 series routers, but this solution uses STP, which is not available on the routers.

Related How To Notes

For a detailed discussion of redundancy options on the server side, see *How to Configure Microsoft Windows 2003 Network Load Balancing Clustering with Allied Telesis Switches*.

The following other How To Notes also use VRRP:

- *How To Configure An L3 Switch To Use Different MAC Addresses On Different VLANs*
- *How To Configure Load Balancer Redundancy on Allied Telesis Routers and Switches*

How To Notes are available from www.alliedtelesis.com/resources/literature/howto.aspx.

Example network

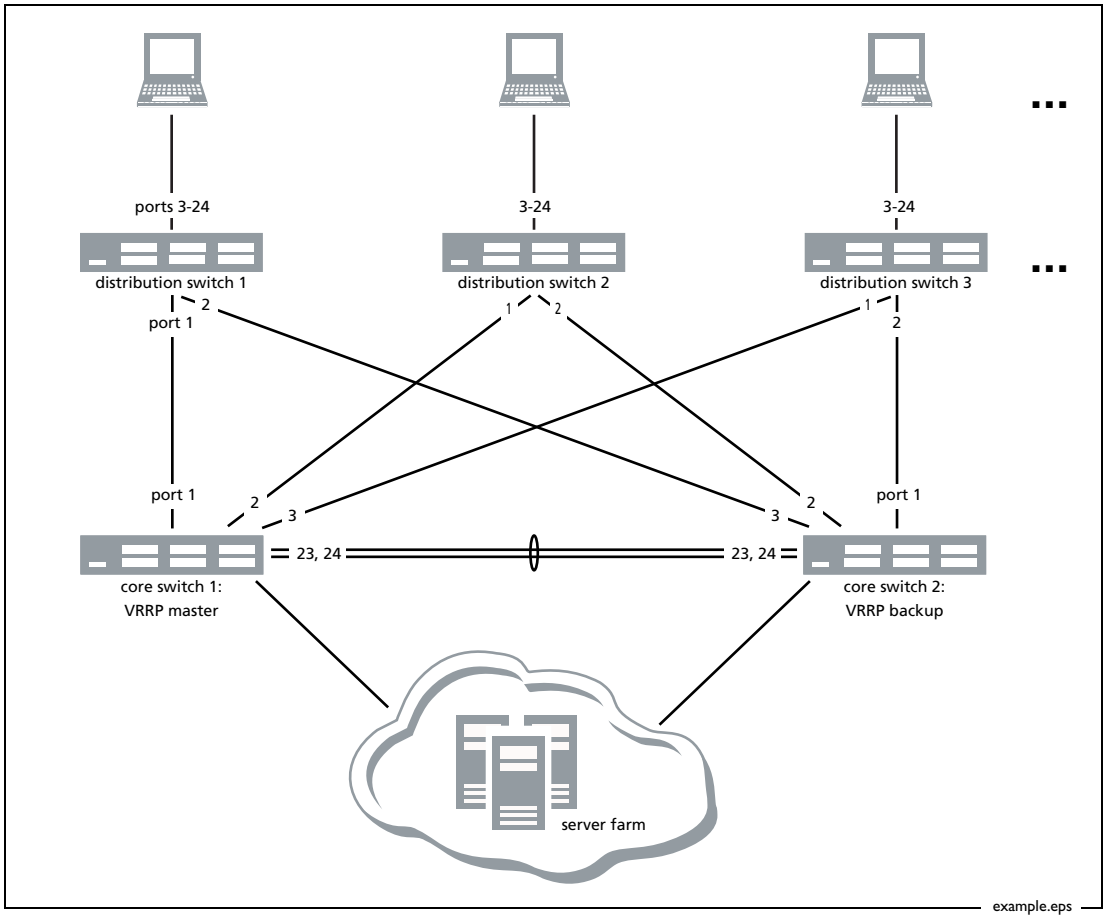

The two core switches are Layer 3 routing and define the IP address ranges for the VLANs configured on the switches. The core switches are also running VRRP to provide router redundancy for end devices such as the PCs and servers. The end devices are connected to the core switches through Layer 2 distribution switches.

The core switches are configured with unique static IP addresses for each VLAN but share a common virtual IP address for each VLAN. VRRP associates a virtual MAC address with

each virtual IP address. The end devices in each VLAN will be configured to use that VLAN's virtual IP address as the default gateway IP address when sending traffic from their local network.

Any end device that wants to send traffic from the local network will ARP for the Layer 2 address of the default gateway. The VRRP master router for the VLAN will respond, sending its VRRP virtual MAC address. Once a response is received from the current master in the virtual router, the end device will send traffic destined for outside the local network to the VRRP virtual MAC address.

Core switch 1 is the VRRP master router for all VLANs, and therefore we configure a higher VRRP priority value on it than on its partner switch core 2. Hence core 1 processes any traffic that has a Layer 2 destination MAC address equal to the virtual MAC address owned by the VRRP master router in any VLAN.

The backup VRRP switch, core 2, shadows the setup of core 1 and is also configured for the same VLANs.

If core 1 fails completely, core 2 will stop receiving VRRP messages from core 1. Core 2 will assume the new forwarding responsibility for all of the VLANs.

If core 1 service is restored, then VRRP message exchanges between core 1 and core 2 will restart, and core 1 will return to the Master state and assume responsibility for all forwarding of traffic sent to the VRRP virtual address. This is the default behaviour and is called *pre-empting*—allowing the VRRP device with the highest priority to assume the master role from another device that is currently acting as the master but that has a lower priority value.

Monitored interfaces

Depending on your server farm configuration, you may need to monitor VLANs on the server side of the core switches. If one of the links from the current VRRP master to the server farm goes down, that switch may need to stop being the master.

To configure this, add the following command to the configuration:

```
add vrrp=<id> monitoredinterface=<vlan>
```

For more information, see the Software Reference and *How to Configure Microsoft Windows 2003 Network Load Balancing Clustering with Allied Telesis Switches*. This How To Note is available from www.alliedtelesis.com/resources/literature/howto.aspx.

Core 1 (master) configuration

1. Name the switch

```
set system name=master
```

2. Create the new VLANs

Create VLANs 2 and 3 for the LANs off distribution switches 2 and 3. The LAN off distribution switch 1 will use VLAN 1.

```
create vlan=vlan2 vid=2
create vlan=vlan3 vid=3
...
```

3. Configure STP

STP is necessary because each distribution switch forms a loop with the two core switches. By default, all VLANs belong to the default STP, so we do not need to add VLANs to the STP.

```
enable stp=default
set stp=default priority=8192 mode=rapid
```

4. Add ports to the VLANs as tagged ports

```
set vlan=1 port=1,23-24 frame=tagged
add vlan=2 port=2,23-24 frame=tagged
add vlan=3 port=3,23-24 frame=tagged
...
```

5. Stop all ports from being untagged ports in VLAN 1

```
delete vlan=1 port=1-24
```

6. Aggregate ports 23-24

```
create switch trunk=example port=23-24
```

The amount of traffic going across the inter-core link depends on the load balancing process being used by the servers. If they use teaming NICs, and some traffic goes via the NIC that is connected to the VRRP backup, then that traffic will go from the server to the backup switch to the master switch to the client (and vice versa). In these circumstances, a lot of traffic could pass between the master and backup switches, so that link needs to be an aggregated set of ports.

7. Give the VLANs IP addresses

```
enable ip
add ip int=vlan1 ip=192.168.1.1
add ip int=vlan2 ip=192.168.2.1
add ip int=vlan3 ip=192.168.3.1
...
```

8. Configure VRRP

```
enable vrrp
create vrrp=1 over=vlan1 ipaddress=192.168.1.3 priority=150
  adoptvrip=on
create vrrp=2 over=vlan2 ipaddress=192.168.2.3 priority=150
  adoptvrip=on
create vrrp=3 over=vlan3 ipaddress=192.168.3.3 priority=150
  adoptvrip=on
...
```

Turning on **adoptvrip** means that the master adopts the VRRP IP address. This means you can use the VRRP IP address to ping the current master, or to connect to it through telnet, SSH, HTTP, SSL or SNMP. Also, DNS relay continues functioning via the same IP address at all times.

Core 2 (backup) configuration

The configuration is very similar to core 1, with the following differences:

- the switch has a different name
- the STP priority value is higher, so core 2 is never the root bridge when core 1 is available
- the VLAN IP addresses are different (but the VRRP virtual IP addresses are the same)
- the VRRP priority is left at the default value of 100.

The complete configuration is:

```
set system name=backup

create vlan=vlan2 vid=2
create vlan=vlan3 vid=3
...

enable stp=default
set stp=default priority=16384 mode=rapid

set vlan=1 port=1,23-24 frame=tagged
add vlan=2 port=2,23-24 frame=tagged
add vlan=3 port=3,23-24 frame=tagged
...

delete vlan=1 port=1-24

create switch trunk=example port=23-24

enable ip
add ip int=vlan1 ip=192.168.1.2
add ip int=vlan2 ip=192.168.2.2
add ip int=vlan3 ip=192.168.3.2
...

enable vrrp
create vrrp=1 over=vlan1 ipaddress=192.168.1.3 adoptvrip=on
create vrrp=2 over=vlan2 ipaddress=192.168.2.3 adoptvrip=on
create vrrp=3 over=vlan3 ipaddress=192.168.3.3 adoptvrip=on
...
```

Distribution switch 1 configuration

The LAN off this switch uses VLAN 1. This exists by default, so you only need to change ports 1 and 2 to tagged ports and enable STP. STP is necessary because each distribution switch forms a loop with the two core switches.

```
set system name=distro1

enable stp=default

set stp=default mode=rapid

set vlan=1 port=1-2 frame=tagged
```

Distribution switch 2 configuration

The LAN off this switch uses VLAN 2, which you need to create.

```
set system name=distro2

create vlan=vlan2 vid=2
add vlan=2 port=3-24
add vlan=2 port=1-2 frame=tagged

enable stp=default
set stp=default mode=rapid
```

Distribution switch 3 configuration

The LAN off this switch uses VLAN 3, which you need to create.

```
set system name=distro3

create vlan=vlan3 vid=3
add vlan=3 port=3-24
add vlan=3 port=1-2 frame=tagged

enable stp=default
set stp=default mode=rapid
```

Testing and troubleshooting the configuration

When a PC attempts to access a server that is behind a VRRP router, the PC starts by ARPing for its gateway IP address, which is the VRRP virtual IP address. The VRRP master then replies with a MAC address of the form 00-00-5e-00-01-xx, where xx is the VRID in hexadecimal (for example 01 for VRRP instance 1).

You can confirm that VRRP is working correctly by looking at the FDB (forwarding database) entries of the distribution switches. The entry for the VRRP virtual MAC address should appear on the port that is connected to the master. For example, if port 1 on the distribution switch is connected to the master switch, you should see the VRRP virtual MAC FDB entry on port 1.

Therefore, to test whether VRRP is working, use the following steps:

1. Enter the command **show vrrp** on each core switch, to check each switch's VRRP status. Core 1 should be the VRRP master and core 2 should be the VRRP backup.

   If the master and backup switches are the wrong way around, check the VRRP priorities.

   If both switches claim to be the master, this is a serious problem, and probably indicates a failure on the communication path between them.

2. From the LAN behind one of the distribution switches, ping one of the servers. The ping should succeed. If it doesn't, check the cabling, the VLAN configuration, and the IP addresses. Also make sure that the PC is using the correct gateway address.

3. On the distribution switch, check that there is an FDB entry for the VRRP MAC address on the port that connects to core 1. This confirms that core 1 replied to the PC's ARP, so core 1 is acting as the master.

4. Power off core 2 (the backup) and ping the server again. The ping should still succeed and the FDB entry should be present for the same port.

   Note that just taking a link down is not an adequate test of this configuration because STP can recover from a broken link. To test VRRP, you need to power the core switch off.

5. Power core 2 back up and ping the server again. There should be no change.

6. Power off core 1, wait a few seconds, and ping the server again. The ping should still succeed because it will travel via core 2. Check the FDB on the distribution switch again. This time, the FDB entry should be on the port that connects to core 2.

   If core 2 doesn't take over as the master when you power core 1 down, check the cabling, the VLAN configuration, and the VRRP configuration.

7. Enter the **show vrrp** command on core 2. It should now be the VRRP master.

If failover happens but is slow, check that network devices are not blocking gratuitous ARPs. The switch sends a gratuitous ARP when it becomes the VRRP master, so that connected switches update their FDB entry for the VRRP virtual MAC. This minimises recovery time because the switches do not have to wait for FDB entries to time out.

USA Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
European Headquarters | Via Motta 24 | 6830 Chiasso | Switzerland | T: +41 91 69769.00 | F: +41 91 69769.11
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
www.alliedtelesis.com

C613-16127-00 REV A

Connecting The IP World

Allied Telesis™