

Understanding and Configuring VIP and Interface Redundancy on the CSS 11000

Document ID: 16556

Introduction

Before You Begin

Conventions

Prerequisites

Components Used

Interface Redundancy

Active/Backup VIP Redundancy

Shared VIP Redundancy

Configurations

Virtual Router Configuration

Interface Redundancy Commands

VIP Redundancy Commands

show Commands

Considerations

Shared VIPs

Combinations

Related Information

Introduction

The purpose of the Virtual IP (VIP) and interface redundancy feature is to provide IP addresses that can float between two or more physical network nodes. These IP addresses are used to provide redundancy for attached servers and VIPs. There are four types of VIP and interface redundancy features: interface, active/backup VIP redundancy, fate sharing, and shared VIP redundancy.

These interfaces perform the following:

- Provide redundancy for interface addresses between switches. The interface address has to be in the range of a subnet common to each of the switches.
- Provide redundancy for VIP addresses between switches. The VIP address has to be in the range of a subnet common to each of the switches. VIP redundancy can be active/backup, where only one switch services requests for the VIP, or shared, where multiple switches service VIP requests.
- Both the IP addresses and MAC addresses of a redundant interface or VIP are shared. In other words, the MAC address does not change when the backup takes over.
- A switch acting as a backup for a particular VIP or interface is up otherwise. This can be contrasted with the box-to-box redundancy feature, where the backup box is considered down, except for the redundancy link.
- In some situations, a box may act as a master for some VIPs and a backup for others.
- VIPs and interfaces on the same subnet can be tied together so that a switch can master a group of VIPs and interfaces (fate sharing). Some topologies require this.
- No special redundancy link between switches is necessary. The redundancy protocol (VRRP) runs over the subnet associated with the redundant interface address or VIP address.

Each redundant interface or VIP is mapped to the combination of an interface on the switch and a Virtual Router ID, or VRID. Each of these combinations defines a virtual router in the terms of the Virtual Router Redundancy Protocol (VRRP). A separate VRRP session is run for each virtual router. Multiple redundant

interfaces or VIPs can be mapped to a single virtual router. If VRRP reports that the switch is master for a VRID on an interface, all redundant interfaces and VIPs associated with the VRID are declared up. Similarly, if VRRP reports that it is in a backup state for a VRID on an interface, all redundant interfaces and VIPs associated with the VRID are declared down. The exception to this is the case of shared redundancy for VIPs. These VIPs are in an active/master state if the switch is master for a VRID and in an active/backup state otherwise.

Only the master for a virtual router will send VRRP messages for that VRID. These messages are MAC-level multicast packets. The source MAC address used for these packets is of the following form:

```
00-00-5E-00-01-{VRID}
```

The use of this address assures that the forwarding databases of any Layer 2 (L2) devices in the network always contain a valid entry for the MAC address associated with the VRID. A redundant interface or VIP master will answer ARP requests for the redundant address. The MAC address in the reply is the one associated with the VRID, as shown above. A switch in backup mode (except for the case of shared VIPs) will forward packets with the VRID's MAC address according to its L2 forwarding table. Content rules that reference the VIP are inactive.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

There are no specific prerequisites for this document.

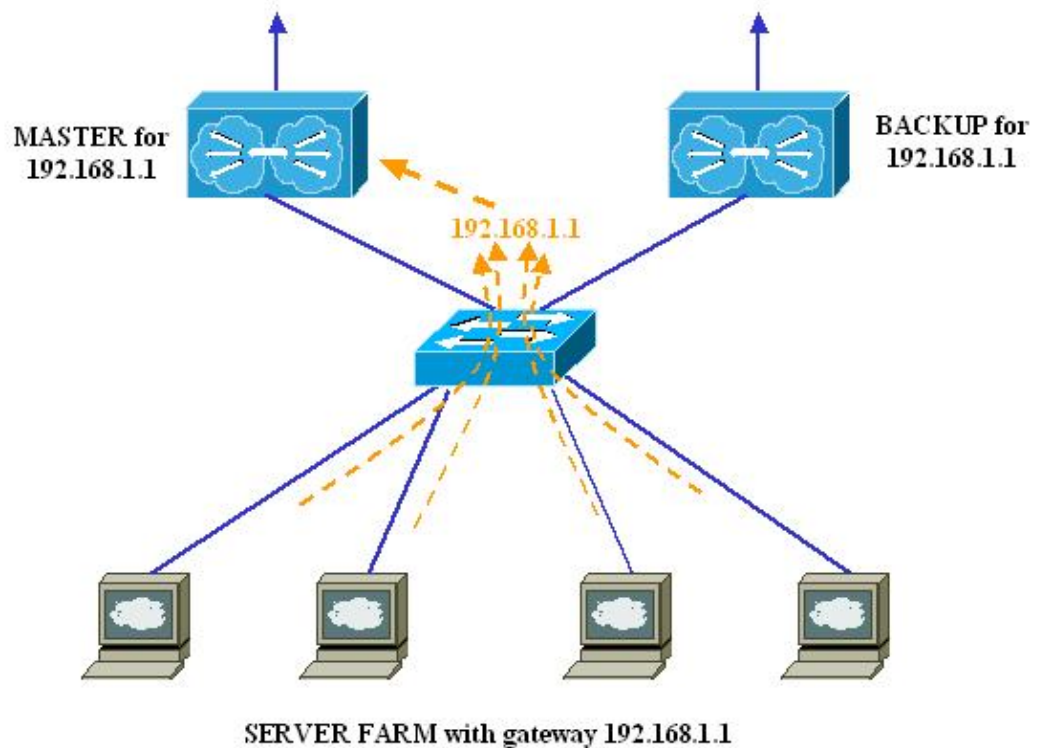
Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Interface Redundancy


CSS 11000 INTERFACE REDUNDANCY



The above diagram shows a number of hosts connected to a L2 switch, which are connected to two Content Services Switch (CSS) 11000 switches. The hosts are configured with a default route that points to 192.168.1.1. Each CSS is configured with a different IP address. Each CSS box is also configured to negotiate mastership for the virtual interface 192.168.1.1. One box will become master for that address. The master box replies to Address Resolution Protocol (ARP) requests and Internet Control Message Protocol (ICMP) echo requests. The hosts will send routed packets to the master. If the master box goes down, one of the backups will assert mastership and send out gratuitous ARPs for the virtual router address in order to update the forwarding tables in the L2 switches (the MAC address does not change).

Active/Backup VIP Redundancy

CSS 11000 ACTIVE/BACKUP VIP REDUNDANCY

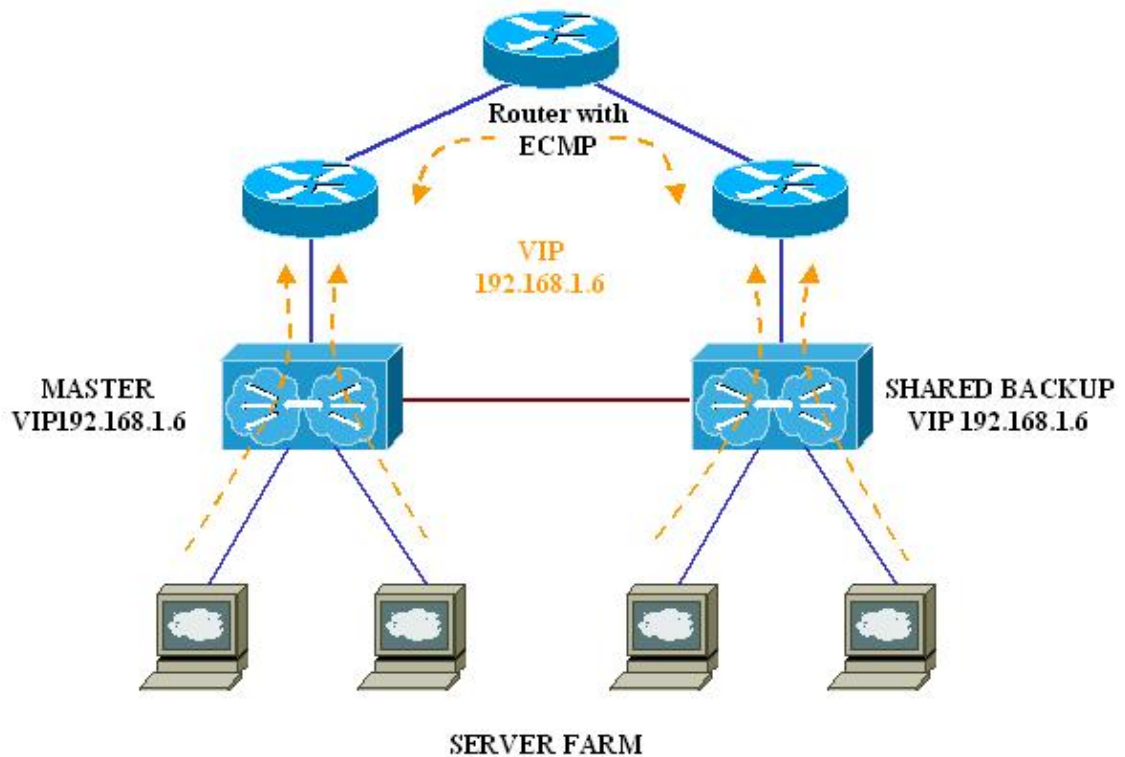


The above diagram shows a number of hosts connected to two CSS 11000 switches. Each CSS is configured with a different IP address. Each CSS is also configured to negotiate mastership for the VIP 192.168.1.5. One box will become master for that address. The master box replies to ARP requests and ICMP echo requests. It also performs the flow setups for flows bound to that VIP. The backup switches treat the VIP address as external to the box. If a backup receives a packet with the VIP as a destination address, it will forward the packet (L2 or L3, as applicable) to the master. All content rules that reference the VIP on the backup switches are inactive.

The routers will send packets whose destination IP address equals the VIP to the master (possibly via one or both backup switches). Router must be directly attached to each CSS.

Shared VIP Redundancy

CSS 11000 SHARED VIP REDUNDANCY



The above diagram is similar to that of active/backup VIP redundancy with some routers on the other side. This time, the switches are configured for shared redundancy for the VIP 192.168.1.6. One box will still become master for that address. The master box replies to ARP requests. All boxes reply to ICMP echo requests and perform flow setups for flows bound to that VIP. This provides VIP load balancing between the switches. The effectiveness of the load balancing depends on the nature of the Equal-Cost MultiPath (ECMP) routing that occurs in the routed network above the switches.

Each router will send packets whose destination IP address equals the VIP to its attached switch, using the MAC address as advertised by the master. When a backup receives such a packet, it will process the packet according to the content rules configured for the VIP and not forward the packet to the master. In order for this to work, there cannot be a shared L2 infrastructure in front of or behind the CSS.

Configurations

The following sections define the new Command Line Interface (CLI) commands used to configure interface and VIP redundancy.

Virtual Router Configuration

Both interface and VIP redundancy are built on top of virtual routers. Virtual router configuration defines the VRID and the parameters used by VRRP.

A virtual router is configured in IP interface mode by issuing the following commands:

```
ip virtual-router vrid [ priority prior ] [ preempt ]  
[ service service name ]  
no ip virtual-router vrid
```

The **vrid field** is an integer between 0 and 255. This number must be unique in the context of the IP interface.

The **priority** keyword is optional. Its argument, **prior**, is an integer between 1 and 255. The switch with the highest priority generally becomes the master for the VRID. The value 255 is known as the owner priority. Only one switch can be configured with the owner priority for a VRID. If the owner of a VRID is alive, it will always assert itself as master. The default priority value is 100.

If two switches have the same priority and (for whatever reason) both have asserted mastership, the one whose (real) interface has the smaller IP address will relinquish mastership to the other node. A backup will never preempt a master with the same priority.

The **preempt** keyword is optional. By default, a switch will not assert itself as master when it is in backup mode and the current master is at a lower priority. If **preempt** is specified, the higher-priority backup will assert mastership in this scenario. This absence or presence of this keyword is ignored if the priority is set to 255. A VRID owner always preempts any other master.

The **service** keyword is optional. If used, the virtual router will only be active when the indicated service is alive. This gives a switch a method to avoid mastering a VRID when it cannot service requests to the associated VIPs.

If the IP interface and any optional service are up and the VRID is new, a VRRP virtual router will be started for the VRID specified. If the VRID matches a virtual router that is already configured, the new parameters will be applied to that virtual router.

The **no** form of the command removes the virtual router.

Interface Redundancy Commands

A redundant interface is configured in IP interface mode.

```
ip redundant-interface vrid ip_address  
  
no ip redundant-interface vrid ip_address
```

The IP address is the address of the redundant interface. This address must meet the following criteria:

- It must be in the range of IP addresses defined by the **ip address** command that put the CLI into interface mode. For example, if the command was **ip address 192.168.1.20/24**, the address must be in the range 192.168.1.1–192.168.1.254, inclusive. The subnet-broadcast addresses (192.168.1.0 and 192.168.1.255) cannot be used.
- It must not match the local IP interface address for that subnet (for example, 192.168.1.20).
- It must not match the IP address of any other redundant interfaces.
- It must not match the IP address of any VIPs.
- It must not match the IP address of any interface on any other node on the subnet.

The **vrid** parameter is an integer between 0 and 255. It must match the VRID of a previous **ip virtual-router** command. All redundant interfaces and VIPs sharing the same fate must use the same VRID. The VRID must be the same value on all CSS switches providing redundancy for this interface.

VIP Redundancy Commands

A redundant VIP is configured in IP interface mode.

```
ip redundant-vip IP address vrid [ shared ]
no ip redundant-vip IP address
```

The IP address is the address of the redundant VIP. This address must meet the following criteria:

- It must be in the range of IP addresses defined by the **ip address** command that put the CLI into interface mode. For example, if the command was **ip address 192.168.1.20/24**, the address must be in the range 192.168.1.1–192.168.1.254, inclusive. The broadcast addresses (192.168.1.0 and 192.168.1.255) cannot be used.
- It must not match the local IP interface address for that subnet (for example, 192.168.1.20).
- It must not match the IP address of any other VIPs.
- It must not match the IP address of any redundant interfaces.
- It must not match the IP address of any interface on any other node on the subnet.

The **vrid** parameter is an integer between 0 and 255. It must match the VRID of a previous **ip virtual-router** command. All redundant interfaces and VIPs sharing the same fate must use the same VRID. The VRID must be the same value on all CSS switches providing redundancy for this VIP.

The optional **shared** parameter is used to indicate shared VIP redundancy. The default is active/backup.

show Commands

The following show commands will be added to the CLI:

```
show virtual-routers
```

Issue the **show virtual-routers [IP interface [vrid]]** command to display state regarding one or more virtual routers.

If no optional parameters are specified, information about all configured virtual routers is displayed. If an IP interface address is included, information about all configured virtual routers on that interface's subnet is displayed. If both an interface address and VRID are specified, information about that single virtual router (if it exists) is displayed.

The display includes the following information:

- The IP interface address.
- The VRID.
- The name of the optional service, if configured.
- The virtual router state.
- If applicable for the virtual router state, the address of the master.
- The amount of time the virtual router has been in the present state.
- The number of state changes that have occurred since the virtual router was configured.

Issue the **show redundant-interfaces [IP address]** to display state regarding one or more redundant interfaces.

If no optional parameter is specified, information about all configured redundant interfaces is displayed. If an IP address that matches an IP interface is included, information about all configured redundant interfaces on that interface's subnet is displayed. If an IP address that matches a redundant interface address is specified, information about that single redundant interface is displayed. If any other IP address is specified, an error message is displayed.

The display includes the following information:

- The redundant interface address.
- The IP interface address.
- The redundant interface's state.
- If applicable for the redundant interface's state, the address of the master.
- The amount of time the redundant interface has been in the present state.
- The number of state changes that have occurred since the redundant interface was configured.

Issue the **show redundant-vips [*IP address*]** command to display state regarding one or more redundant VIPs.

If no optional parameters are specified, information about all configured redundant VIPs is displayed. If an IP address that matches an IP interface is included, information about all configured redundant VIPs on that interface's subnet is displayed. If an IP address that matches a redundant VIP address is specified, information about that single redundant VIP is displayed. If any other IP address is specified, an error message is displayed.

The display includes the following information:

- The redundant VIP address.
- The IP interface address.
- The redundant VIP's state.
- If applicable for the redundant VIP's state, the address of the master.
- The amount of time the redundant VIP has been in the present state.

The number of state changes that have occurred since the redundant VIP was configured.

Considerations

Shared VIPs

When shared VIP redundancy is configured, it is very important that no devices exist between the CSS boxes and the routers. This is important for the following reasons:

- If a L2 network interconnects the CSS boxes and the routers, all traffic will go to the CSS that is master/active, thus defeating the purpose of shared redundancy.
- If a L2 switch or hub divides a router and a CSS, the router will not see its link go down upon CSS failure. The router will continue to advertise the subnet containing the VIP as reachable. The shared redundancy feature counts on the router to stop advertising the subnet if the CSS or the link between the CSS and the router goes down.

The bottom line is that each router must connect to exactly one CSS.

On the server side, the reverse flow problem is difficult to solve if there is a shared L2 infrastructure. Instead, each server should be isolated to a single CSS. Each CSS should be linked to all others providing redundancy for the same VIPs.

Combinations

The desired configuration may consist of a CSS acting as master for one set of VIPs, backup for another set, and possibly sharing for yet another collection of VIPs. If that is the case, it is imperative that the return traffic from the server hits a CSS before any L2 switch in order to have the reverse NAT transformation performed.

Related Information

- [Application Networking Services Support Resources](#)
 - [Technical Support – Cisco Systems](#)
-

All contents are Copyright © 2006–2007 Cisco Systems, Inc. All rights reserved. [Important Notices and Privacy Statement](#).

Updated: Aug 03, 2007

Document ID: 16556
