VRRP allows for transparent failover at the first-hop IP router by configuring a group of routers to share a VIP address. VRRP selects a primary router in that group to manage all packets for the VIP address. The remaining routers are in standby and take over if the primary router fails.

A LAN client can determine which router should be the first hop to a particular remote destination by using a dynamic process or static configuration. Examples of dynamic router discovery are as follows:

Proxy ARP: The client uses ARP to get the destination it wants to reach, and a router responds to the ARP request with its own MAC address.

Routing protocol: The client listens to dynamic routing protocol updates and forms its own routing table.

ICMP Router Discovery Protocol (IRDP) client: The client runs an ICMP router discovery client.

The disadvantage of dynamic discovery protocols is that they incur some configuration and processing overhead on the LAN client. Also, if a router fails, the process of switching to another router can be slow.

An alternative to dynamic discovery protocols is to statically configure a default router on the client. Although this approach simplifies client configuration and processing, it creates a single point of failure. If the default gateway fails, the LAN client is limited to communicating only on the local IP network segment and is cut off from the rest of the network.

VRRP can solve the static configuration problem by enabling a group of routers (a VRRP group) to share a single VIP address. You can then configure the LAN clients with the VIP address as their default gateway.

In this example in the figure, Routers A, B, and C form a VRRP group. The IP address of the group is the same address that was configured for the Ethernet interface of Router A (10.0.0.1). Because the VIP address uses the IP address of the physical Ethernet interface of Router A, Router A is the primary router (also known as the IP address owner). As the primary router, Router A owns the VIP address of the VRRP group and forwards packets that are sent to this IP address. Clients 1 through 3 are configured with the default gateway IP address of 10.0.0.1.

Routers B and C function as backups. If the primary router fails, the backup router with the highest priority becomes the primary router and takes over the VIP address to provide uninterrupted service for the LAN hosts. When router A recovers, it becomes the primary router again.

The benefits of VRRP are as follows:

Redundancy: This enables you to configure multiple routers as the default gateway router, which reduces the possibility of a single point of failure in a network.

Load sharing: This benefit allows multiple routers to share traffic to and from LAN clients. The traffic load is shared more equitably among available routers.

Multiple VRRP groups: This benefit supports up to 255 VRRP groups on a router physical interface if the platform supports multiple MAC addresses. Multiple VRRP groups enable you to implement redundancy and load sharing in your LAN topology.

Multiple IP addresses: This benefit allows you to manage multiple IP addresses, including secondary IP addresses. If you have multiple subnets that are configured on an Ethernet interface, you can configure VRRP on each subnet.

Pre-emption: This benefit enables you to pre-empt a backup router that has taken over for a failing primary router with a higher-priority backup router that has become available.

Advertisement protocol: This benefit uses a dedicated Internet Assigned Numbers Authority (IANA) standard multicast address (224.0.0.18) for VRRP advertisements. This addressing scheme minimizes the number of routers that must service the multicasts and allows test equipment to accurately identify VRRP packets on a segment. IANA has assigned the IP protocol number 112 to VRRP.

VRRP tracking: This benefit ensures that the best VRRP router is the primary router for the group by altering VRRP priorities based on interface states.

The VRRP primary router sends VRRP advertisements to other VRRP routers in the same group. The advertisements communicate the priority and state of the primary router. Cisco NX-OS encapsulates the VRRP advertisements in IP packets and sends them to the IP multicast address 224.0.0.18 that is assigned to the VRRP group. Cisco NX-OS sends advertisements once every second by default, but you can configure a different advertisement interval.

VRRP supports the following two options for tracking:

Native interface tracking: This option tracks the state of an interface and uses that state to determine the priority of the VRRP router in a VRRP group. The tracked state is down if the interface is down or if the interface does not have a primary IP address.

Object tracking: This option tracks the state of a configured object and uses that state to determine the priority of the VRRP router in a VRRP group.

If the tracked state (interface or object) goes down, VRRP updates the priority based on what you configure the new priority to be for the tracked state. When the tracked state comes up, VRRP restores the original priority for the virtual router group. For example, you may want to lower the priority of a VRRP group member if its uplink to the network goes down so another group member can take over as primary router for the VRRP group.

VRRP supports high availability through stateful restarts and Stateful Switchovers. A stateful restart occurs when the VRRP process fails and is restarted. Stateful Switchover occurs when the active supervisor switches to the standby supervisor. Cisco NX-OS applies the run-time configuration after the switchover.

Multiple VRRP Groups
You can configure up to 255 VRRP groups on a physical interface. The number of VRRP groups that a router interface can support depends on the following factors:

Router processing capability

Router memory capability

In a topology where multiple VRRP groups are configured on a router interface, the interface can act as a primary router for one VRRP group and as a backup for one or more other VRRP groups.

This topology contains two VIP addresses for two VRRP groups that overlap. For VRRP group 1, Router A is the owner of IP address 10.0.0.1 and is the primary router. Router B is the backup to Router A. Clients 1 and 2 are configured with the default gateway IP address of 10.0.0.1.

For VRRP group 2, Router B is the owner of IP address 10.0.0.2 and is the primary router. Router A is the backup to router B. Clients 3 and 4 are configured with the default gateway IP address of 10.0.0.2.

VRRP Router Priority and Pre-Emption
An important aspect of the VRRP redundancy scheme is the VRRP router priority, because the priority determines the role that each VRRP router plays and what happens if the primary router fails. If a VRRP router owns the VIP address

and the IP address of the physical interface, this router functions as the primary router. The priority of the primary router is 255.

The priority also determines if a VRRP router functions as a backup router and the order of ascendancy to becoming a primary router if the primary router fails. For example, if Router A, the primary router in a LAN topology, fails, VRRP must determine if backups B or C should take over. If you configure Router B with priority 101 and Router C with the default priority of 100, VRRP selects Router B to become the primary router because it has the higher priority. If you configure Routers B and C with the default priority of 100, VRRP selects the backup with the higher IP address to become the primary router.

VRRP uses pre-emption to determine what happens after a VRRP backup router becomes the primary router. With pre-emption enabled by default, VRRP switches to a backup if that backup comes online with a priority higher than the new primary router. For example, if Router A is the primary router and fails, VRRP selects Router B (next in order of priority). If Router C comes online with a higher priority than Router B, VRRP selects Router C as the new primary router, even though Router B has not failed.

If you disable pre-emption, VRRP switches only if the original primary router recovers or the new one fails.

VRRPv3 and Virtual Router Redundancy Service
VRRP version 3 (VRRPv3) enables a group of switches to form a single virtual switch to provide redundancy and reduce the possibility of a single point of failure in a network. The LAN clients can then be configured with the virtual switch as their default gateway. The virtual switch, representing a group of switches, is also known as a VRRPv3 group.

Virtual Router Redundancy Service (VRRS) improves the scalability of VRRPv3 by providing a stateless redundancy service to VRRS pathways and VRRS clients by monitoring VRRPv3. VRRPv3 acts as a VRRS server that pushes VRRPv3 status information (such as current and previous redundancy states, active and inactive Layer 2 and Layer 3 addresses, and so on) to VRRS pathways and all registered VRRS clients.

VRRS clients are other Cisco processes or applications that use VRRPv3 to provide or withhold a service or resource dependent upon the state of the group. VRRS pathways are special VRRS clients that use the VRRS database information to provide scaled first-hop gateway redundancy across scaled interface environments.

VRRS by itself is limited to maintaining its own state. Linking a VRRS client to a VRRPv3 group provides a mechanism that allows VRRS to provide a service to client applications so that they can implement stateless or Stateful Failovers. A Stateful Failover requires communication with a nominated backup before the failure so that operational data is not lost when the failover occurs.

VRRS pathways operate in a similar way to clients but are integrated with the VRRS architecture. They provide a means to scale first-hop gateway redundancy by allowing you to configure a virtual address across hundreds of interfaces. The virtual gateway state of a VRRS pathway follows the state of an FHRP VRRS server.

VRRPv3 notifies VRRS of its current state (primary, backup, or nonoperational initial state [INIT]) and passes that information to pathways or clients. The VRRPv3 group name activates VRRS and associates the VRRPv3 group with any clients or pathways that are configured as part of VRRS with the same name.

Pathways and clients act on the VRRPv3 server state. When a VRRPv3 group changes states, VRRS pathways and clients alter their behavior (performing tasks such as shutting down interfaces or appending accounting logs) depending on the state that is received from VRRS.