;login:

THE MAGAZINE OF USENIX & SAGE

December 2001 • Volume 26 • Number 8

inside:

SECURITY

High Availability Firewall/VPN with VRRP

By Dave Zwieback

USENIX & SAGE

The Advanced Computing Systems Association &
The System Administrators Guild

high availability firewall/VPN with VRRP

by Dave Zwieback

Dave Zwieback is the technical director of inkcom (*http://www.inkcom.com*), a company that specializes in system, network, and security architecture.

zwieback@inkcom.com

Introduction

Internet connectivity has become mission-critical for many organizations, especially as they shift their connections to customers and partners from private lines to virtual private networks (VPNs). Not surprisingly, the availability requirements for Firewalls and VPNs have substantially increased. In this article, I discuss the implementation of a high availability (HA) Firewall/VPN using the Virtual Router Redundancy Protocol Monitored Circuit (VRRPmc). Specifically, I cover Check Point Firewall-1/VPN-1 version 4.1 SP3 running on appliances made by Nokia, with IPSO version 3.3. However, since most network equipment manufacturers currently support VRRP version 2 (RFC 2338) and VRRPmc, the general concepts discussed here apply to almost any HA network configuration.

Virtual Router Redundancy Protocol Monitored Circuit

If static routing is used and a router fails, users have to manually change their configuration to point to a replacement router. Using dynamic routing protocols (RIP, OSPF, BGP, etc.) allows for route replacement to happen automatically after a timeout. While dynamic routing clearly provides higher availability than static routing, dynamic configurations are more difficult to manage and can result in significant network overhead. In addition, transition from a failed router may take an unacceptably long time, a condition known as "black hole" periods.

VRRP is designed to combine the simplicity of static routing with the high-availability features of dynamic routing. In a VRRP Monitored Circuit configuration, users point to a static IP address of a virtual router. This virtual router has valid IP and MAC addresses, which under normal conditions are "owned" by the master router (i.e., the master forwards packets sent to the IP address of the virtual router and responds to appropriate ARP requests). In the event of master router failure, a standby backup router becomes master and assumes ownership of the virtual router (see Diagram 1). Typically, users experience no downtime during failure.

Each virtual router must have a Virtual Router Identifier (VRID) in the range of 1 to 255. Although each virtual router may have more than one IP address, all the IP addresses of a particular virtual router are associated with one MAC address (from the address block assigned by IANA specifically for VRRP), and the VRID is the last octet: 00:00:5E:00:01:VRID. Thus, while two virtual routers with the same VRID can successfully exist on different LANs, VRIDs must be unique on a particular LAN to prevent MAC address conflicts.

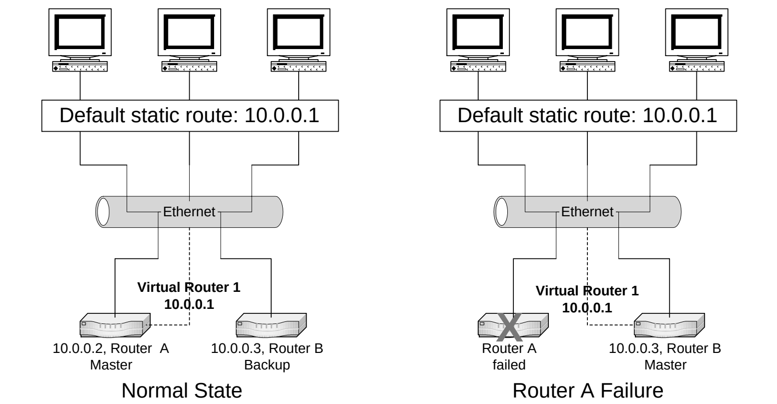

Diagram 1

A physical VRRP router may participate in (or "back up") more than one virtual router. For instance, a router may be master for one VRID and backup for another – a typical "active-active" configuration, illustrated in Diagram 2. This scenario is useful for load balancing between two (or more) routers, and offers increased performance in addition to high availability.

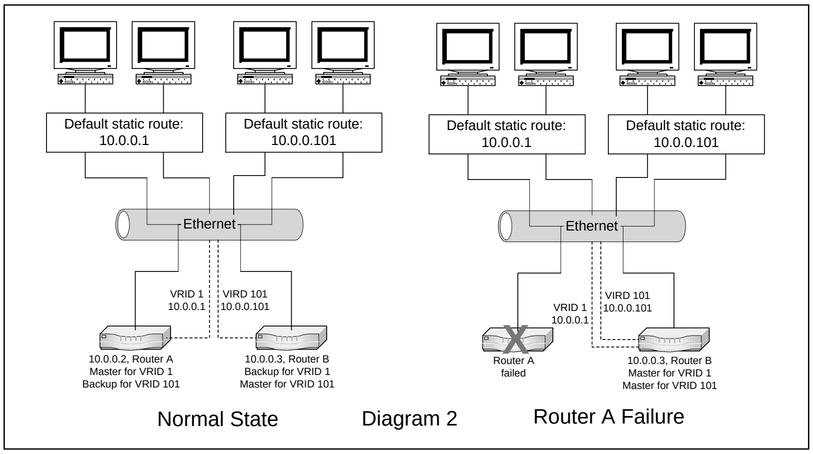

Diagram 2

To determine which VRRP router is master or backup at any given time, each router must be assigned a priority value between 1 and 255 for each VRID. The router with the highest priority is the master. The master sends periodic advertisement messages to a special VRRP multicast address, 224.0.0.18. The frequency of these messages is typically 1 second, which can be changed by adjusting the "Hello Interval." When the master stops sending messages (for instance, in case of complete failure like a power outage), the backup router with the highest priority will take over the virtual router after a brief (<< 1 second) timeout.

For example, in an active-active configuration, detailed in Diagram 2, priorities can set be as follows:

	VRID 1	**VRID 101**
Router A	100	95
Router B	95	100

In this setup, Router A is master for VRID 1 and backup for VRID 101; Router B is master for VRID 101 and backup for VRID 1. Should Router A fail, VRID 1 is transferred to Router B, because it has the highest priority (95) of all routers participating in VRID 1 at that point. Since, Router B still has the highest priority (100) for VRID 101, the failure of Router A does not have any effect on VRID 101.

If a router with a higher priority than the current master comes online (for instance, a failed master router is fixed and becomes available), the virtual router moves to the router with the highest priority. Once again, users should not experience any downtime.

One of the most important differences between VRRPmc and VRRP (version 2) is how interface failure is treated. In VRRPv2, when an interface fails, only the failed one is transferred to the backup. On the other hand, in a VRRPmc configuration, the router can be configured to constantly monitor its physical interfaces (thus "monitored circuit"), and in case any of them fail, *all* the affected VRIDs are transferred to the backup. VRRPmc prevents asymmetric routing, which is required for the proper operation of the Check Point Firewall-1/VPN-1 in an HA configuration.

With either VRRPv2 or VRRPmc, the failover is accomplished by reducing the priority for the affected VRIDs by a specified amount (known as the "Priority Delta"). Once the priority of the master is lower than that of any of the backup routers, those with the highest priority immediately take over the VRIDs. Once again, the users should not experience a significant connectivity outage.

On a final note, VRRP is supposed to provide for strong authentication between participating routers, but currently the only choices on the Nokia platform are either no authentication or simple text passwords. It is highly recommended that the latter be used, especially in potentially "hostile" firewall environments.

Configuring VRRPmc on Nokia Appliances

The following is a sample Nokia appliance VRRPmc configuration, as detailed in Diagram 3. There are three networks: External, Internal, and DMZ. The two firewalls, A

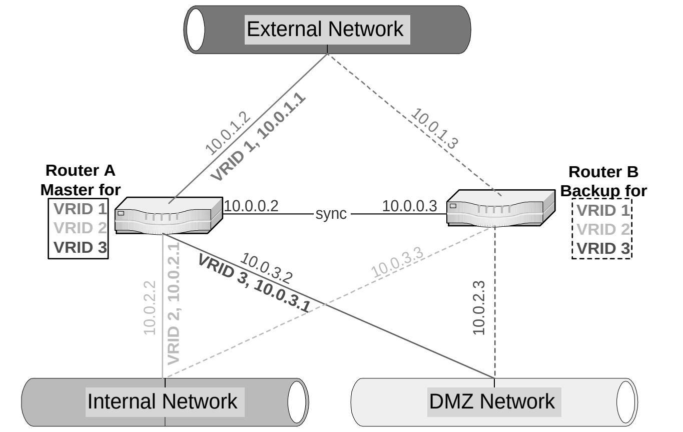

Diagram 3

and B, are connected to all three networks, and are also connected to each other via a crossover cable for Firewall-1 state synchronization (discussed below). Under normal conditions, Firewall A will serve as master for the three VRIDs (one on each network), while Firewall B will serve as backup.

The VRRPmc configuration can be accessed through the Voyager interface by going to Config→Router Services→VRRP menu. The VRRP settings for this configuration are summarized in the following table. Note that all the VRIDs are set up in monitored-circuit mode.

	Router A Master	Router B Backup
External Interface	10.0.1.2	10.0.1.3
Virtual Router (VRID)	1	
Priority	100	95
Hello Interval	1	
Backup IP (VRRP IP)	10.0.1.1	
Monitor Interfaces	Internal, Priority Delta: 10 DMZ, Priority Delta: 10	
Authentication	Simple Password: ViP1	
Internal Interface	10.0.2.2	10.0.2.3
Virtual Router (VRID)	2	
Priority	100	95
Hello Interval	1	
Backup IP (VRRP IP)	10.0.2.1	
Monitor Interfaces	External, Priority Delta: 10 DMZ, Priority Delta: 10	
Authentication	Simple Password: ViP2	
DMZ Interface	10.0.3.2	10.0.3.3
Virtual Router (VRID)	3	
Priority	100	95
Hello Interval	1	
Backup IP (VRRP IP)	10.0.3.1	
Monitor Interfaces	Internal, Priority Delta: 10 External, Priority Delta: 10	
Authentication	Simple Password: ViP3	

Each interface is configured to monitor two other interfaces (the only interface that is not monitored is the one used for firewall state synchronization). Thus, in case any of the monitored interfaces fails on Firewall A, the priority of all the VRIDs will be reduced to 90 (original priority of 100 minus Priority Delta of 10). Since the priority of the backup for these VRIDs is 95, all the VRIDs will be immediately transferred to Firewall B, which will become the new master. If Firewall A comes back online with priority higher than 95, it will take over the VRIDs once again.

Configuring HA Firewall-1/VPN-1 with Gateway Clusters

There are three steps involved in configuring the Check Point Firewall-1/VPN-1 for high availability:

1. Configuring the state synchronization
2. Configuring Gateway Clusters
3. Allowing VRRP traffic in the rule base

Step 1

Check Point Firewall-1/VPN-1 keeps track of all connections in its state table. In order to facilitate failover between two (or more) firewalls, this state information needs to be synchronized. On each firewall, the $FWDIR/conf/sync.conf file must contain a list (IP addresses or resolvable names) of all the firewalls that the state table should be synchronized with. In addition, a "control path" must be established between the firewalls, using the fw putkey command. Furthermore, since state information is exchanged approximately every 100 milliseconds, it is recommended that a separate network interface be dedicated to the task on each firewall and that time is synchronized between the firewalls as well, for instance, using xntp. A typical configuration is presented in Diagram 3.

With the firewall states synchronized, when the master fails, connections originally passing through the master continue through the backup uninterrupted. ($FWDIR/lib/table.def can be modified to include or exclude specific tables or protocols from state synchronization). One visible difference during the failover period is in the firewall log, where the *origin* of log entries will now be the backup instead of the original master firewall.

Step 2

With the introduction of Gateway Clusters, Check Point has considerably simplified HA firewall and VPN configuration. A Gateway Cluster is a virtual firewall, which consists of two or more physical firewalls configured for HA, for instance with VRRPmc and state synchronization. The steps involved in configuring Gateway Clusters are as follows:

1. Verify that the management console is separate from any of the HA firewalls (otherwise Gateway Clusters will not work).
2. Check "Enable Gateway Clusters" in the Policy➔Properties➔High Availability tab.
3. Create a Gateway Cluster object, with the external VRRP IP address: Manage➔Network Objects➔New➔Gateway Cluster.
4. Modify the properties of each of the member firewalls to make them members of the Gateway Cluster created in step 3.

Once the firewalls become members of the Gateway Cluster, their individual Authentication, VPN, and Certificate settings disappear from their properties and can now be modified through the Gateway Cluster object. Furthermore, the security policy is now installed on the Cluster instead of on the individual member firewalls. However, by default, the policy install will fail unless it is successful on all the members of the Cluster. (This can be changed by checking off Policy→Properties→Security Policy→Install Security Policy only if it can be successfully installed on ALL selected targets, and by checking off Policy→Properties→High Availability→Install Security Policy on Gateway Cluster only if it can be successfully installed on ALL Gateway Cluster members.)

Traffic from the Gateway Cluster may be coming from either the "real" or the VRRP IP addresses. It is important to take this into account when creating firewall rules and objects. Specifically, in VPN configurations, if the remote firewall is using Gateway Clusters and is managed from a different management console, the remote firewall object should be created with the Cluster address and have all the "real" IP addresses for all the firewalls in the Cluster specified in the Interfaces Tab.

STEP 3

As mentioned before, the master router periodically sends VRRP advertisements to 224.0.0.18. The firewalls should have a rule allowing VRRP advertisements:

1. Verify that the VRRP service is defined (with Match: ip_p=0x70).
2. Create the VRRP-MCAST-NET workstation object with an address of 224.0.0.18.
3. Create a rule in the security policy to allow VRRP traffic as follows:

Source	Destination	Service	Action	Track	Install On
MasterFW BackupFW	VRRP-MCAST-NET	VRRP	Accept	Long	FWCluster

As mentioned before, if you use the "Install On" field, ensure that Gateway Cluster object (FWCluster in the example above) is used instead of individual firewalls.

Monitoring and Troubleshooting

The VRRP status can be viewed from the Nokia Voyager, as well as from the command line, using iclid. Following are the iclid commands related to VRRP:

show vrrp	quick summary report
show vrrp interface	interface configuration
show vrrp stat	vrrp stats

Furthermore, on the current master, ifconfig displays the VRID IP and MAC addresses along with the actual interface information.

```
FirewallA# ifconfig -a
…
inet 10.0.1.1/24 broadcast 10.0.1.255 vrrpmac 0:0:5e:0:1:1
inet 10.0.1.2/24 broadcast 10.0.1.255
…
```

By default, a VRRPmc IP address cannot be pinged, although this functionality can be enabled in IPSO 3.3. In addition to being useful for troubleshooting purposes, enabling this feature is required for certain routers and operating systems that will not forward any traffic to a gateway that does not respond to pings (for instance, the "dead gateway detection" in HPUX). Because it is not always possible to know all the specific

Using VRRP provides an easy way to greatly improve availability and performance of the network.

devices that utilize a particular router, it is recommended to enable "Accept Connections to VRRP IP" in the Voyager VRRP menu, and drop or reject *and* log icmp traffic with a firewall rule; this way, it is possible to identify devices that are pinging the router, which is extremely useful in troubleshooting.

Aside from configuration errors, most problems with VRRP installations occur when the backup router stops receiving VRRP messages from the master. This causes the backup to assume that the original master has failed, and to transfer to become master. However, if the original master did not actually fail and is still in master state, this "split-brain" situation may cause all sorts of havoc on the network. This condition is most often caused by VRRP traffic (to 224.0.0.18) being blocked, for instance by a firewall rule (or lack thereof) or an incorrect router or switch ACL or VLAN configuration. Please refer to Nokia Support Resolution 1521 for an excellent checklist of typical VRRP problems and solutions.

Furthermore, there are known issues with some Ethernet switches, mostly due to the fact that when a failover situation occurs, the MAC address (00:00:5E:00:01:VRID) of the virtual router "moves" from one switch port to another. The most common symptom is excessive VRRP transition master-to-backup time (sometimes over 30 seconds!), and thus a connectivity outage. Where appropriate, hubs can be used instead of switches, especially if the backup router is only used in standby mode. Otherwise, disabling MAC address caching and the spanning tree algorithm on the appropriate ports is required (on Cisco switches, set port fast will also work).

Conclusion

Using VRRP provides an easy way to greatly improve availability and performance of the network. Specifically, Nokia appliances running Check Point Firewall-1/VPN-1 can be easily configured for high availability with VRRPmc and Gateway Clusters. In addition to transparent failover in case of malfunction of one of the firewalls, the benefits of such a configuration include reduced maintenance and increased performance (in an active-active setup).

Bibliography

Check Point Software Technologies Ltd. 2000. Check Point VPN-1/FireWall-1 Administration Guide.

Knight et al., April 1998. Virtual Router Redundancy Protocol, Request for Comment 2338: *http://www.ietf.org/rfc/rfc2338.txt*.

Nokia Support: *https://support.nokia.com/*.

PhoneBoy's FireWall-1 FAQ: *http://www.phoneboy.com/*.