



From IP-Replacement to IP-Aliasing: Handling Version changes in HACMP

*By Mark Neuman
Technical Services Project Manager
Datatrend Technologies*

In the latest version of IBM's HACMP, version 5.1, there has been a shift in the way HACMP handles service IPs. They went from IP-replacement to IP-aliasing; therefore, both the system administrator and the networking group must take this into account. A service IP is the IP address which all of the clients of a HACMP protected system use to contact the application. It is considered a resource that will follow the applications in the event of a system takeover, thus allowing the actual physical server of the applications to be transparent to any clients.

This change showed up in the version 4.5 of HACMP, as an additional method of managing the service IP. IP-aliasing has become the preferred method of managing service IPs since it is much faster than IP-replacement and can cut the recovery time of an adapter failure by more than half. In addition, it allows you to support more than one service IP on a system without the corresponding increase in the number of network adapter cards needed, as several service IPs can be supported on a single adapter if desired.

In the past, a HACMP protected server would have at least 2 network adapters, one adapter, the service adapter, would have a "boot" IP assigned to it that would be in existence when the HACMP subsystems were NOT operational and would host the service IP once HACMP was started. By the way, both of these IP's needed to be on the same logical subnet. The other adapters would be "standby" adapters, which need to be on a separate logical subnet, although in the same VLAN or physical connection.

Starting with HACMP Version 5.1, the standard method of supporting the service IP is to use IP-aliasing, although IP-replacement is still supported and can be configured. In IP-aliasing there are no longer any "standby" adapters, they are all considered "non-service" adapters. Each adapter MUST be on a separate logical subnet. This may cause problems with your networking group. They almost always need education about how HACMP works and complain about your using all of their IPs and subnets. For example, all of the adapters and their subnets MUST be in the same VLAN along with the service IP network range. This is contrary to what a CISCO trained network engineer will want to do, so be prepared to explain the logic of IP takeover to them.

An additional requirement to using IP-aliasing for the service IPs is that it is only available to networks that support gratuitous ARP. In the event of a takeover, you must deal with ARP issues as hardware address takeover is NOT available. One side effect of using IP-aliasing for your service IP is that you will have an IP (the non-service IP) that you can use to maintain the system, and start/stop HACMP that will not change underneath you, causing you to disconnect just at the time you want to maintain visibility to the process changes. There are many additional changes to HACMP, version 5.1, than the addition of IP-aliasing to handle service IP's, but this will be the most visible shift in the way HACMP works. The HACMP administrator, however, will see a significant change as all of the SMIT screens have changed, along with the methodology of setting up a HACMP cluster.

***Datatrend's TrendSetter eNewsletter
July 15, 2004***