

# CLUVFY's Complaint & “10.”

*Prutser Certified Paper #1 version 1.0*

*Harald van Breederode. (Harald.van.Breederode@Oracle.com)*

*Principal DBA Instructor, Oracle University NL*

## 1 Introduction

During the installation of the Oracle 10g release Clusterware (CRS) a new tool is used to verify the cluster setup. This Cluster Verification Utility or CLUVFY (pronounced *cluffy*) can be used by the DBA before he or she starts the Oracle Universal Installer (OUI) but is also executed by the OUI after the CRS is installed to verify its success. One of the things CLUVFY checks is that the IP addresses assigned by you or your network administrator to the RAC Virtual IP (VIP) are usable based on certain criteria and if not then CLUVFY returns errors.

In certain situations CLUVFY will indeed complain about the VIP's and the aim of this paper is to discuss this issue and understand whether or not this is a problem. This in turn requires some basic knowledge about the IP protocol, IP addresses, netmasks and subnetting. If you have this knowledge you can skip directly to the conclusion section but if so then perhaps you wouldn't be reading this paper!! So lets start with the basics.

## 2 The IP protocol

The Internet Protocol (IP) was developed in the early 1980's as a protocol to allow computers to communicate with each other over a network. Despite its name, it was not designed for the Internet as we know it today. It was designed for DARPA net which was a network for the US military. Later the IP protocol was used to build the ARPA network for research institutes, universities and other non-commercial companies. And finally this ARPA network became The Internet and was made available to the rest of the world and became the Internet as we know it today. If we speak about an Internet also known as an intranet, it means an IP based network and if we speak about “The Internet” it means what your mother knows as internet. The IP protocol used today is known as IP version 4 or Ipv4. In order to communicate over an Internet each system must have a unique address, called an IP address.

## 2.1 The IP address

An IP address consists of 4 bytes or octets  $X1.X2.X3.X4$  and is usually written as decimal numbers, eg 101.102.103.104. This IP address can be broken down into a network number and a host number within that network. A network packet or datagram can be directly sent from one IP address to another IP address if both systems are on the same network, but if they don't share the same network number then the packets must be sent to a router which knows the route from one network number to the other. Thus a router sends packets from one network to another network. In order for a router to successfully route a packet it must quickly divide the IP address into a network number and a host number at a known boundary within the four bytes of the IP Address.

## 2.2 Network Classes

To support routing in different sized networks the IP protocol caters for large, medium and small networks using network classes which divide the total IP address space into Classes A,B,C,D and E networks as follows:

### 2.2.1 Class A networks

Class A networks use IP addresses where the first bit is always 0. Therefore Class A IP addresses have an  $X1$  value between 1 and 127 inclusive denoting the network number, whilst  $X2,X3$  and  $X4$  denote the host number within that network. Therefore there are 127 possible Class A networks of roughly 16 million hosts each.

**Note:** Network 127.X.X.X is reserved for loopback purposes.

### 2.2.2 Class B networks

Class B networks use IP addresses where the first 2 bits are always 1 and 0 (10). Therefore class B IP addresses have an  $X1$  value between 128 and 191 inclusive which together with  $X2$  octet denote the network number, whilst  $X3$  and  $X4$  denote the host number within that network. Therefore there are roughly 16000 possible Class B networks of roughly 64000 hosts each.

### 2.2.3 Class C networks

Class C networks use IP addresses where the first 3 bits are always 1, 1 and 0 (110). Therefore class C IP addresses have an  $X1$  value between 192 and 223 inclusive which together with  $X2$  and  $X3$  octets denote the network number whilst  $X4$  denotes the host number within that network. Therefore there are roughly 2 million possible Class C networks of 256 hosts each.

### 2.2.4 Class D and E networks

IP addresses above 224.x.x.x are reserved for multicasting purposes among other things and are known as Class D and Class E addresses.

## 2.3 Coping with the shortage of IP network numbers

When the Internet was new there was no shortage of IP network numbers available but the Internet Engineering Task Force (IETF) realised that due to the growth in Internet usage, IP network numbers would soon be exhausted requiring a solution. The first step was restricting the allocation of registered network numbers to computers directly connected to the Internet and the assignment of special ranges of numbers in each network class for systems not directly connected to the Internet. These freely available network numbers were declared non-routable meaning they are not allowed on the Internet but can be routed within your intranet. Transmission of IP packets with a non-routable IP address over The Internet would get you fired by the IP Police Department (IETF). These reserved network number ranges are: 10.X.X.X for Class A, 172.16.X.X to 172.31.X.X for Class B and 192.168.X.X to 192.168.255.X for Class C and are available for any purpose providing they don't get transmitted on the Internet.

The second step was the introduction of the Network Address Translation Protocol also known as NAT allowing the router to translate reserved network numbers inside a private network into a registered IP address when transmitting packets to the Internet and reversing the process for incoming packets. With NAT you could use a reserved network number inside your network and still be connected to the Internet using only one registered IP address thereby greatly reducing the need for registered IP addresses. Today NAT is widely used, enabling the Internet to grow more than would have been possible without it. Eventually however, even NAT will not suffice but it has given the IETF time to design the definitive solution called Ipv6 which is not considered in this paper.

## 2.4 The Netmask

A system may send an IP packet to another system directly over the network if both systems share the same network number. If not, then the packet needs to be sent to a router first which knows where to route the packet on another network. The network and a host numbers in an IP address are determined by the use of the netmask which specifies which bits contain the network number and which bits contain the host number within that network. Each address class has its own netmask. For Class A addresses the default netmask is 255.0.0.0 meaning that all bits of X1 are used for the network number. For Class B addresses the default netmask is 255.255.0.0 meaning that all bits of X1 and X2 are used for the network number and for Class C addresses the default netmask is 255.255.255.0 meaning that all bits of X1, X2 and X3 are used for the network number. Performing a logical *AND* operation on the binary values of the IP address and the netmask results in the network number. All bits that are not part of the network number based on the masking bits form the host number within that network.

## 2.5 Subnetting and Supernetting

Occasionally Class A, B or C networks are too large for requirements. Sometimes Class B or C networks are too small, or a Class B or C network is too small and a Class A or B is too large. Supernetting and Subnetting permit the combination of smaller networks into larger ones or the subdivision of larger ones into smaller ones respectively, by extending or shortening the usual netmask for the class. For example by using a Class B network with address 140.85.0.0 and a netmask of 255.255.255.0 instead of the default 255.255.0.0, the Class B network is broken or subnetted into 256 networks of 256 hosts instead of 16,000 networks of 65,536 hosts by adding extra bits to the network number portion of the address and removing bits from the host number portion. Alternatively, to combine Class C networks 200.200.200.0 and 200.200.201.0 into one larger network of 512 hosts instead of 2 networks of 256 hosts each, we would use a netmask of 255.255.254.0 rather than 255.255.255.0 which removes 1 bit from the network number portion of the IP address and adds 1 bit to the host number portion thereby doubling the number of possible hosts in the network at the cost of reducing the number of networks.

**2.5.1:** In the mid 90's routers began to use Classless Internet Domain Routing (CIDR) and the concept of IP address classes began to fade away. CIDR uses a different notation for the netmask. Instead using four decimal numbers CIDR only encodes the numbers of 1's in the netmask. So a netmask of 255.0.0.0 is written as /8. Usually you see something like 140.85.240.0/20 specifying a network number 140.85.240.0 with a netmask of 255.255.240.0 thus a subnetted class B network.

## 3 RAC VIP's and Reserved Networks and Subnet Masks

Now that we have a basic understanding of the IP protocol, IP addresses, netmasks, subnetting and supernetting we can focus on the question of CLUVFY and the issue of what may be wrong with using VIP addresses of 10.X.X.X or any other non-routable IP address? The answer is "it depends"!!

If you use network 10.X.X.X for your company IP network and your RAC cluster is isolated from the Internet, then there is no problem. But if your RAC cluster must be reachable over the Internet via your router then this will result in errors because packets for network 10.X.X.X or any other non-routable IP address cannot be sent over the Internet. CLUVFY has been designed to complain if you try to use an IP address which belongs to a non-routable and reserved network number in any Class because these were originally declared non-routable even though that that is technically not true within an intranet. It doesn't matter whether or not your reserved network number is subnetted; CLUFVY will complain. I believe CLUVFY may be interpreting this issue too strictly and if your cluster need NOT be reachable from the Internet then you can ignore CLUVFY's complaint.

### 3.1 VIPCA May Also Complain

In addition to the CLUVFY issues raised in this paper the Virtual IP Address Configuration Assistant utility (VIPCA) seems to use the wrong netmask by default and if not overridden with the correct subnet mask, then connection problems will occur regardless of which IP address class is used.

## **4 Conclusion**

There is nothing wrong with a VIP address of 10.X.X.X or other so called non-routable address as long as your RAC cluster is isolated from The Internet. CLUVFY is too strict in it's VIP address checks. In other words – CLUFVY complains too loudly.

### ***About the author***

I work for Oracle University, The Netherlands as Principal DBA Instructor teaching most of the Oracle DBA curriculum courses with specialisation in High-Availability, Data Warehousing and New Features courses as well as doing technical review for DBA curriculum development. In addition I have a strong background in Unix System Administration and Networking having worked 6 years for Oracle as a Unix System Administrator in the Dutch Data Center before switching to Oracle University. I have Oracle 8i, 9i and 10g Oracle Certified Professional and Oracle 10g Certified Master DBA qualifications and am a member of the Oak Table Network.

In addition I specialise in accessibility issues for the visually impaired having lost my sight a few years ago, thereby necessitating my use of computer accessibility hardware and software for work and home. I use JAWS as my screen reader in combination with an ALVA Satellite refreshable Braille display write JAWS scripts as required to make inaccessible programs accessible to me.

### **Final Note**

*I wish to thank my dear friend Joel Goodman for raising this question, encouraging me to write this paper and for reviewing, correcting and formatting it.*