

VirtualCluster

Virtual Clusters

Hadoop Clusters can be created on Virtual Machines. See [Virtual Hadoop](#) for the benefits and challenges present in such deployment.

Process to set up a cluster using the Hadoop shell scripts

1. All machines should be brought up to date, with the same version of Java, Hadoop, etc.
2. All machine's(both VM's and physical machines) public key are distributed to all "~/.ssh/authorized_keys" file.
3. conf/hadoop-site.xml file is similar for all the machines.
4. /etc/hosts file must contain all the machines(VM,Physical machine) IP and Hostname.
5. The local hostname entry in /etc/hosts must not point to 127.0.0.1 or any other loopback address (some laptop-friendly Linux distributions do this). It should be to the assigned IP address.
6. conf/slaves must contain the hostname of all slaves including VM's and physical machine.
7. conf/masters must contain only master's hostname.
8. both conf/masters and conf/slaves files must be similar in all the participating machines.
9. The hadoop home directory should be same for all the machines like "/root/hadoop-0.18.2/"
10. Finally execute the Namenode -format command only in the Master machine.

Most of this configuration can be done on a single machine image, the problematic area is host naming. Virtual networks with reverseDNS allow the hostname to be picked up from the network. If you use this,

1. The namenode and job tracker hostnames should remain constant; these need separate images.
2. Bring up other machines afterwards; use an empty conf/slaves file.

Trouble Areas

Here are things that can cause trouble.

1. Multiple virtual network adapters. It is simpler with one network adapter/node
2. Machines changing hostname/IPAddress on a reboot. For a long-lived virtual cluster you need stable machine names.
3. Machines whose hostname doesn't match the hostname the network assigns it. It thinks it is "granton", the network thinks it is "dhcp-169-45", that being the name everything else talks to it by.
4. Machines that think they have the same hostname. You get this if you clone VMs and don't rename them.
5. Pauses of an entire VM for 5-10s or longer. This happens when the virtual host is overloaded and your VM has been swapped out. Host less VMs, or have them ask for less memory.
6. Wierd clock drift where it can even run backwards. Again, don't overload your machines.
7. All redundant virtual servers (e.g. Namenode and secondary NN) being hosted on the same physical machine. At that point, you don't have redundancy or failover any more.