Cray @ Duke: Cluster Install Instructions

Alainna White (awhite@cray.com)
12/18/2013

Summary

This document describes how to install CentOS 6.4 on cluster nodes in the Duke AMD Gigabit Cluster. A CentOS 6.4 x64 Linux install has been provided. Cray has generated an RSA SSH Key for the cluster (betanew:/root/.ssh/id_rsa[.pub]). The public key is included in the kickstart for passphrase-less access to the nodes from the head node.

This document is provided as a courtesy; there is no implied warranty for this configuration as it was done pro bono.

1. INSTALLING NEW NODES

This section describes how to install (AKA "Kickstart") nodes that are new to the cluster. This should include the previously obtained AMD nodes.

1.1 CONFIGURE NODES FOR PXEBOOT

The nodes will need to be configured to first look for an install pointer on the TFTP server. This is achieved by entering the BIOS of the target host. These instructions are for the Supermicro H8QGL-IF motherboard.

Steps:

1. Reboot the node, and press DEL to run Setup.
2. Go to the "Boot" tab in the BIOS
3. Select "Boot Device Priority". Make sure Network: "IBA GE Slot" is the first device and the hard drive is the second device.
4. Press F10 to save your configuration.
5. Reboot the node.

6. The machine will try to PXEboot. During this phase a MAC/Hardware address is presented. Write it down. You'll need it for the next stage. (If you already know the MAC/Hardware addresses of the Ethernet ports then you can skip this step).

1.2 ADD NEW HOSTS TO HEAD NODE

There are two places that the hosts will need to be added to the head node to do a network install: The `/etc/hosts` file and the `/etc/dhcp/dhcpd.conf` file. **IMPORTANT: Before making a change to the `dhcpd.conf` file, make sure you create a backup copy first.** If you edit the `dhcpd.conf` file and the syntax is incorrect, DHCPD will not restart until you fix the error. It's safer to revert to the previous version of the file, restart DHCPD and then redo your edits and try again.

/etc/hosts

The syntax for adding new hosts to the hosts file is [IP ADDRESS] [HOSTNAME]. For example:

```
10.10.1.1       prod-0001
10.10.1.2       prod-0002
…
10.10.1.123     prod-0123
10.10.1.xxx     prod-xxx
```

Add the host's desired IP and hostname and save the file.

/etc/dhcp/dhcpd.conf

The entire group that is managed by the DHCP server obtains values from the server that tells it where and what to pull from the TFTP server in order to PXEBoot. This next block just for reference only; you won't have to modify this unless you change servers.

```
[Snip]
allow booting;
allow bootp;
next-server 10.10.1.254;
filename "pxelinux.0";
option space pxelinux;
option pxelinux.magic code 208 = string;
option pxelinux.configfile code 209 = text;
option pxelinux.pathprefix code 210 = text;
option pxelinux.reboottime code 211 = unsigned integer 32;

[Snip]
        class "pxeclients" {
        next-server 10.10.1.254;
        filename "pxelinux.0";
```

You will have to add the new host to DHCP. You'll need the MAC (Hardware) address for the first (eth0) network adapter that you obtained from the BIOS.

Go to the bottom of the list of hosts and add your new hosts before the last "}".
The syntax for adding a new host to the dhcpd configuration file is as follows:

```
host prod-0123 { hardware ethernet 00:25:90:5A:36:56; fixed-address 10.10.1.123; }
```

*** BE SURE that your syntax is correct! If it isn't, DHCPD will fail on reload!**

Once you have added your hosts, you can issue the a "service dhcpd restart" and "service dhcpd status" command.

```
[root@betanew ~]# service dhcpd restart
Shutting down dhcpd:                                       [  OK  ]
Starting dhcpd:                                            [  OK  ]
[root@betanew ~]# service dhcpd status
dhcpd (pid  43116) is running...
[root@betanew ~]#
```

1.3 SETUP THE INSTALLER (HEAD NODE)

The next step is to let the head node know which nodes you want to install.

On the head node, enter the `/tftpboot/linux-install/pxelinux.cfg` directory.

This is what the file listing looks like for the `pxelinux.cfg` directory:

```
z[root@betanew pxelinux.cfg]# ls -l
total 20
drwxr-xr-x 2 root root 4096 Dec 18 10:39 bak
-rw-r--r-- 1 root root  169 Dec 17 14:19 centos6
-rw-r--r-- 1 root root   72 Dec 17 15:04 default
-rwxr-xr-x 1 root root  254 Dec 17 14:18 iptohex
-rw-r--r-- 1 root root  279 Dec 11 17:03 pxeos.xml
```

The network boot agent will look for installation pointer files with a series of names that pertain to the hexadecimal version of its IP Address; if it cannot locate a relevant pointer file, it will default to the file named `default`. The `default` file tells the system to continue on with a local boot. **It is important that no hexadecimal files exist if no installs are running**. If a hexadecimal file exists for a node or subset of nodes, it will be picked up by the nodes' network boot instead of the default file. Most likely this will result in the node(s) not booting to their local operating system.

The `centos6` file is the configuration pointer that tells the node where to pull the installation from and other important directions for the install.

To initiate a kickstart for a node:

1. Use the `iptohex` script to generate the hexadecimal version of the node(s) IP address:

```
[root@betanew pxelinux.cfg]# ./iptohex 10.10.1.104
0A0A0168
```

2. Next, you want to create a symbolic link between the centos6 file and the hexadecimal address of your node(s):

```
[root@betanew pxelinux.cfg]# ln -s centos6 0A0A0168
[root@betanew pxelinux.cfg]# ls -l 0A0A0168
lrwxrwxrwx 1 root root 7 Dec 18 11:07 0A0A0168 -> centos6
```

3. Reboot the node(s).

4. Once the installer is running on the node, delete the pointer file. Failure to remove the pointer file will result in an install loop.

```
[root@betanew pxelinux.cfg]# rm -rf 0A0A0168
```

5. Once the node has finished the installation and rebooted, you can SSH from the head node to make sure the install was successful and connectivity/SSH passphrases were properly configured.

6. Enjoy your newly-installed OS!

2. REINSTALLING EXISTING NODES

This section describes how to reinstall nodes. This may be necessary after a hard drive failure, as a debug measure, etc.

In these instances, the server and PXEBoot client are already configured. Therefore, configuring the head node to serve a hexadecimal pointer file to the relevant host is all that needs to be done.

2.1 SETUP THE INSTALLER (HEAD NODE)

Let the head node which nodes you want to install.

On the head node, enter the `/tftpboot/linux-install/pxelinux.cfg` directory.

This is what the file listing looks like for the `pxelinux.cfg` directory:

```
[root@betanew pxelinux.cfg]# ls -l
total 20
drwxr-xr-x 2 root root 4096 Dec 18 10:39 bak
-rw-r--r-- 1 root root  169 Dec 17 14:19 centos6
-rw-r--r-- 1 root root   72 Dec 17 15:04 default
-rwxr-xr-x 1 root root  254 Dec 17 14:18 iptohex
-rw-r--r-- 1 root root  279 Dec 11 17:03 pxeos.xml
```

The network boot agent will look for installation pointer files with a series of names that pertain to the hexadecimal version of its IP Address; if it cannot locate a relevant pointer file, it

will default to the file named `default`. The `default` file tells the system to continue on with a local boot. **It is important that no hexadecimal files exist if no installs are running**. If a hexadecimal file exists for a node or subset of nodes, it will be picked up by the nodes' network boot instead of the default file. Most likely this will result in the node(s) not booting to their local operating system.

The `centos6` file is the configuration pointer that tells the node where to pull the installation from and other important directions for the install.

To initiate a kickstart for a node:

7. Use the `iptohex` script to generate the hexadecimal version of the node(s) IP address:

```
[root@betanew pxelinux.cfg]# ./iptohex 10.10.1.104
0A0A0168
```

8. Next, you want to create a symbolic link between the centos6 file and the hexadecimal address of your node(s):

```
[root@betanew pxelinux.cfg]# ln -s centos6 0A0A0168
[root@betanew pxelinux.cfg]# ls -l 0A0A0168
lrwxrwxrwx 1 root root 7 Dec 18 11:07 0A0A0168 -> centos6
```

9. Reboot the node(s).

10. Once the installer is running on the node, delete the pointer file. Failure to remove the pointer file will result in an install loop.

```
[root@betanew pxelinux.cfg]# rm –rf 0A0A0168
```

11. Once the node has finished the installation and rebooted, you can SSH from the head node to make sure the install was successful and connectivity/SSH passphrases were properly configured.

12. Enjoy your newly-reinstalled OS!